

PhysMAP - interpretable *in vivo* neuronal cell type identification using multi-modal analysis of electrophysiological data

Eric Kenji Lee¹, Asim Emre Gül², Gregory Heller⁸, Anna Lakunina^{4,5,6}, Santiago Jaramillo^{4,5}, Pawel F. Przytycki⁹, Chandramouli Chandrasekaran^{1,3,7,10,*}

¹Department of Psychological and Brain Sciences, Boston University, Boston, 02115, MA, USA

²Department of Psychology, Boğaziçi University, Beşiktaş, 34342, İstanbul, Türkiye

³Department of Anatomy & Neurobiology, Boston University, Boston, 02118, MA, USA

⁴Department of Biology, University of Oregon, Eugene, 97403, OR, USA

⁵Institute of Neuroscience, University of Oregon, Eugene, 97403, OR, USA

⁶Electrophysiology Group, Allen Institute for Neural Dynamics, Seattle, 98109, WA, USA

⁷Department of Biomedical Engineering, Boston University, Boston, 02115, MA, USA

⁸Department of Brain and Cognitive Science, Massachusetts Institute of Technology, Cambridge, 02139, MA, USA

⁹Faculty of Computing & Data Sciences, Boston University, Boston, 02215, MA, USA

¹⁰Center for Systems Neuroscience, Boston University, Boston, 02115, MA, USA

Abstract

Cells of different types perform diverse computations and coordinate their activity during sensation, perception, and action. While electrophysiological approaches can measure the activity of many neurons simultaneously, assigning cell type labels to these neurons is an open problem. Here, we develop PhysMAP, a framework that weighs multiple electrophysiological modalities simultaneously in an unsupervised manner and obtain an interpretable representation that separates neurons by cell type. PhysMAP is superior to any single electrophysiological modality in identifying neuronal cell types such as excitatory pyramidal, PV⁺ interneurons, and SOM⁺ interneurons with high confidence in both juxtacellular and extracellular recordings and from multiple areas of the mouse brain. PhysMAP built on ground truth data can be used for classifying cell types in new and existing electrophysiological datasets, and thus facilitate simultaneous assessment of the coordinated dynamics of multiple neuronal cell types during behavior.

Keywords: Cell types, Electrophysiology, Opto-tagging, Multi-Modal Analysis, Dimensionality Reduction

(Abstract: 137 words)

1 Introduction

2 Single cell transcriptomics, *in vivo* and *in vitro* electrophysiology, and morphological reconstruction
3 have identified numerous neuronal cell types in the cerebral cortex (Scala et al., 2021; Gouwens et al.,
4 2020; , BICCN) each with specific patterns of gene expression, morphology, and connectivity (Bomkamp
5 et al., 2019). These interconnected cell types form cortical microcircuits which perform neural computa-
6 tions and, ultimately, produce behavior (Keller et al., 2020; Kullander and Topolnik, 2021; Christensen
7 et al., 2022). For instance, recent studies have identified how interactions between excitatory pyramidal,
8 parvalbumin-positive (PV⁺), and somatostatin-positive (SOM⁺) cells enable gain modulation (Ferguson

*Corresponding Author (cchandr1@bu.edu)

9 and Cardin, 2020). These interactions have been used to model computation for functions such as the
10 selective enhancement of certain visual stimuli (Millman & Ocker et al. 2020).

11 However, how individual cell types contribute to the many other computations the brain conducts is
12 largely unknown (Adesnik and Naka, 2018), as direct measurement of specific types necessitates molecular
13 techniques (e.g., optogenetics and calcium imaging) that are not feasible in all situations. Such techniques
14 are far from turnkey and are not easy to execute due to various factors e.g., a lack of viral serotype efficacy
15 (Bohlen and Tremblay, 2023), the need for optical access and head-fixation (O’Shea et al., 2017), or
16 the depth of relevant neural populations (Trautmann et al., 2021) etc. Even when these techniques are
17 available, technical constraints restrict experimentation: only a few cell types are imageable at once and
18 two-photon imaging is limited to cortical depths of around 600 μm (Takasaki et al., 2020). This precludes
19 the study of cell types from deep structures and in many of the species most relevant for understanding
20 human disorders such as non-human primates (Bliss-Moreau et al., 2022). Furthermore, in order to
21 develop precision circuit-level therapies, cell type-specific activity is needed to validate microcircuit models
22 as network properties are predictive of therapeutic response (Bloch et al., 2022). Thus, our current
23 inability to capture cell type dynamics hinders our understanding of how circuits orchestrate behavior in
24 intact and diseased states.

25 In contrast to optical methods, extracellular electrophysiological recordings from single electrodes to
26 high-density Neuropixels probes (Jun et al., 2017) can sample from diverse and deep populations of
27 neurons simultaneously and require no genetic or optical access. This recording technique is also easily
28 scalable to many cortical areas simultaneously (Steinmetz et al., 2019; Siegle & Jia et al. 2021; Chen
29 & Liu et al. 2024). However, in these experiments, the cellular identities of the recorded neurons are
30 almost entirely invisible. While atlases for identifying cell types from their transcriptomic expression are
31 available (Tasic et al., 2018; Bakken et al., 2021), equivalent “cell type classifiers” that can match cell
32 type to specific *in vivo* electrophysiological properties do not exist. This is because the correspondence
33 between cell types and their *in vivo* electrophysiological properties are not understood.

34 Classically, one solution was to use one-dimensional measures such as the action potential width to
35 differentiate “broad-spiking” putatively excitatory from “narrow-spiking” putatively fast-spiking inhibitory
36 neurons (Mitchell et al., 2007; Hussar and Pasternak, 2009). However, such a division neglects differences
37 in waveform shape important for differentiating cell types (Lee et al., 2021; Amatrudo et al., 2012;
38 Povysheva et al., 2013; Krimer et al., 2005; Zaitsev et al., 2008) which is ultimately the result of
39 transcriptomic differences in the expression of ion channels (Bomkamp et al., 2019; Bakken et al., 2021).
40 Another notable electrophysiological property that differs between cell types is a neuron’s intrinsic firing
41 rate pattern called its inter-spike interval (ISI) distribution (Latuske et al., 2015; Schneider et al., 2023).
42 In addition, the circuit connectivity of various cell types is often distinct leading to differences in the input
43 that a cell receives (Zaitsev et al., 2008). These differences in input, in combination with a neuron’s
44 intrinsic dynamics, shape firing rate pattern in responses to external stimuli and task events also known
45 as a peri-stimulus time histogram (PSTH) (Pinto and Dan, 2015; Yu et al., 2019; Ramadan et al., 2022).
46 Thus, each modality provides different but complementary information about a cell and therefore an ideal
47 approach would be to combine these different modalities according to their reliability or informativeness in
48 delineating differences in cell type. However, it’s unclear which of these “electrophysiological modalities”
49 are most important or else how they can be combined to create cell type classifiers that allow for
50 the monitoring of multiple cell types *in vivo*. Therefore, optimally combining the information present
51 across different modalities is critical for building electrophysiological cell type classifiers and ultimately,
52 uncovering the function of neural circuits.

53 Here we present “PhysMAP”, an approach that reliably identifies different cell types by multi-modal inte-

54 gration of different physiological modalities such as a neuron’s waveform shape, ISI distribution, PSTH,
55 derived electrophysiological metrics, etc. PhysMAP achieves multi-modal combination by combining non-
56 linear dimensionality reduction (McInnes et al., 2018) with a weighted-nearest neighbor graph (WNN)
57 construction from Seurat v4, a tool that has seen considerable success in multiomics (Hao & Hao et al.
58 2021). PhysMAP offers three key features: First, PhysMAP can incorporate any number of electrophys-
59 iological properties and arrive at a representation that captures underlying cell types better than any one
60 modality (such as waveform shape) alone. Second, structure elicited by PhysMAP separates neurons
61 by cell type and can delineate certain cell types in a manner that we demonstrate adheres closely to
62 high-level “ground truth” labels available through optogenetic tagging (Yu et al., 2019; Lakunina et al.,
63 2020; Petersen et al., 2021). Third, PhysMAP has the ability to learn a representation and infer cell
64 types in new, unlabeled data allowing it to function as a cell type classifier. Thus, PhysMAP unlocks
65 the study of circuit dynamics of multiple neuronal cell types simultaneously during sensation, perception,
66 and action.

67 (Introduction: 769 words)

68 Results

69 *Multi-modal analysis of extracellular electrophysiology with PhysMAP*

70 Our PhysMAP approach uses a weighted graph combination solution developed in multiomics for intelligently
71 integrating multiple modalities (transcriptomic, epigenomic, and proteomic data) in order to find cell type struc-
72 ture that transcends single modalities (Hao et al., 2021). As in our previous WaveMAP approach (Lee et al.,
73 2021, 2023), PhysMAP uses the average normalized data for each single unit. Unlike WaveMAP, PhysMAP looks
74 at arbitrarily many modalities beyond simply waveform shape. PhysMAP then finds shared underlying structure
75 by combining these modalities via a weighted nearest neighbor (WNN) graph. After assessing all modalities for
76 all cells, this method unifies all modality-specific graphs weighing each according to their informativeness on a
77 per-unit basis. This graph is then projected into two dimensions for visualization of high-dimensional multi-modal
78 structure. In Fig. 1, we show how this WNN graph is constructed in the simple two-modality case combining
79 waveform shape and ISI distribution. What follows here is a description of these steps appended by the motivation
80 for each. These steps are also described in greater detail in [Methods: Weighted Nearest Neighbors](#).

- 81 1. **Within-Modal Affinity:** Within the waveform space (Fig. 1A), we select a single neuron (blue sphere
82 in Fig. 1B), identify its k -nearest neighbors (in this example, $k=5$; red spheres in Fig. 1B), and average
83 them to predict the waveform of said neuron (dashed red circle in Fig. 1C). We then calculate the *within-*
84 *modal affinity* by passing the neuron’s actual waveform and its predicted waveform into a modified UMAP
85 distance kernel (Fig. 1D, numerator), and thus a measure of how well a neuron’s waveform is predicted by
86 its neighbors in waveform-space. In the third step, we will show how this will be useful when determining
87 which modalities to “trust” more in the form of an “affinity ratio”.
- 88 2. **Cross-Modal Affinity:** Next, we calculate a *cross-modal affinity* for the ISI distribution modality (Fig.
89 1E). Using the same neuron and its same neighbors in Fig. 1B, but now in ISI-space where each dimension
90 is a time point along an ISI distribution curve (Fig. 1F), we calculate a predicted ISI distribution for this
91 neuron by averaging (dashed red circle in Fig. 1G). We do this again by passing the neuron’s true ISI
92 distribution and its predicted ISI distribution into a modified UMAP distance kernel (denominator of Fig.
93 1D), providing a measure of how well one modality versus another is predictive of a neuron’s properties
94 using the same neighbors as to facilitate comparison.
- 95 3. **Affinity Ratio:** We define the per-unit *waveform affinity ratio*, $S_{\text{Wave}}(i)$, (fraction in Fig. 1D) as the ratio
96 of this neuron’s waveform *within-modal affinity* divided by its ISI dist. *cross-modal affinity*. Conversely, the
97 ISI distribution per-unit affinity ratio $S_{\text{ISI}}(i)$ is also obtained (not shown) by considering the ISI distribution-
98 space as the within-modality and the waveform shape-space as the cross-modality. This ratio compares
99 how well one modality versus another is predictive of a neuron’s properties and will be used to determine
100 how much to weigh each modality for this neuron.

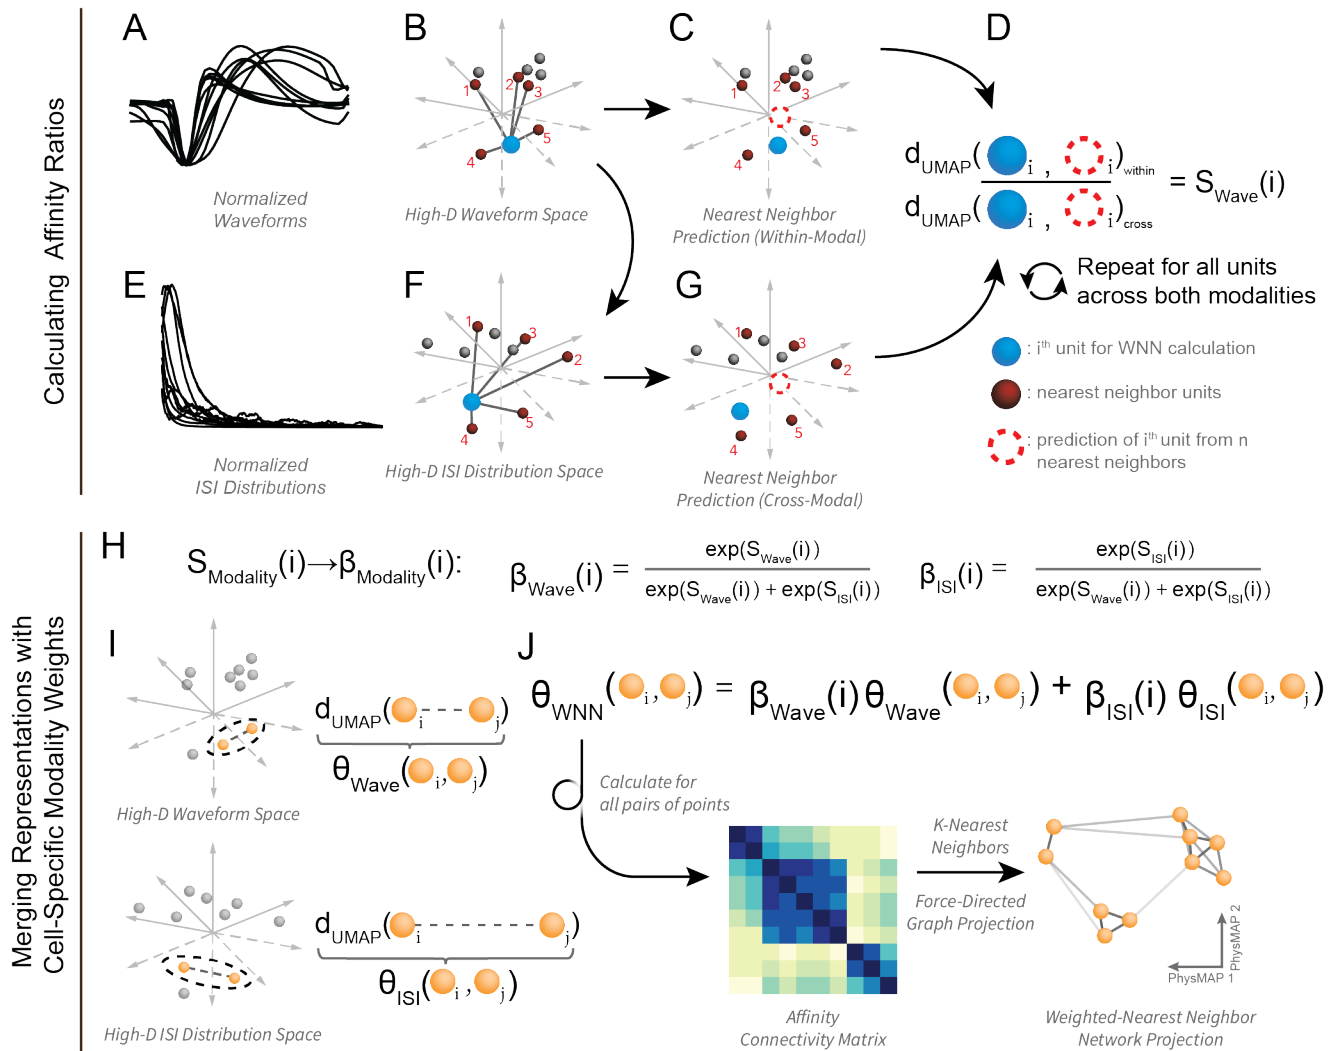


Figure 1: Schematic of PhysMAP and weighted-nearest neighbors algorithm. (A) A sample of normalized average single unit extracellular action potential waveforms. (B) Waveforms in (A) are shown in high-dimensional space with each axis pertaining to each time point along the waveform’s trace. For a sample unit (blue sphere), its nearest neighbor waveforms are highlighted (numbered red spheres). (C) The nearest neighbor waveforms are averaged to generate a prediction of the original waveform (dashed red circle). (D) The numerator of the waveform affinity ratio ($S_{\text{Wave}}(i)$), the “within-modal affinity”, is taken as the difference between the original neuron’s average waveform and its nearest neighbor prediction. This “reconstruction error” distance is passed through a modified UMAP distance kernel. (E) Normalized ISI distributions drawn from the same neurons that produced the waveforms in (A). (F) The same neuron in (B; blue sphere) and its same nearest neighbors in waveform space are now shown in ISI distribution space. (G) As in (C), a prediction is made for the sample unit but this time in ISI-space (red dashed circle). The difference between these, passed again through a modified UMAP distance kernel, is used to form the “cross-modal affinity” that appears in the denominator of the affinity ratio in (D). The ratio of within- and cross-modal affinities form the waveform affinity ratio ($S_{\text{Wave}}(i)$) for the sample unit. This ratio is calculated for each unit’s waveform and ISI affinity ratios ($S_{\text{Wave}}(i)$ and $S_{\text{ISI}}(i)$). (H) Per-unit affinity ratios ($S_{\text{Modality}}(i)$) for each modality are converted into per-unit modality weights ($\beta_{\text{Wave}}(i)$ and $\beta_{\text{ISI}}(i)$) via these formulae. (I, top) The difference between a sample pair of points in high-dimensional waveform space is another type of “affinity” ($\theta_{\text{Wave}}(i, j)$). This is found by passing the distance into UMAP’s distance kernel as before. (I, bottom) The same pair of points as in (I, top) with the UMAP distance between them in high-dimensional ISI space ($\theta_{\text{ISI}}(i, j)$). (J) The equation for calculating weighted affinities (θ_{WNN}) as a linear combination used in constructing the weighted-nearest neighbors graph. This weighted average of affinities is calculated for all points to construct an affinity connectivity matrix. The final multi-modal projection is found after applying k -nearest neighbors to the affinity connectivity matrix and projecting into two dimensions with UMAP’s projection algorithm (a force-directed graph layout procedure).

- 101 4. **Modality Weight:** We then calculate the weight associated with a unit for a particular modality ($\beta_{\text{Modality}(i)}$)
102 as the ratio of the exponentiated affinity ratio for a modality divided by the sum of exponentiated affinity
103 ratios over all modalities (Fig. 1H; left equation for waveforms, right equation for ISI dists.). Thus for each
104 neuron, each modality is differentially weighted according to how well the cell's properties are predicted in
105 each.
- 106 5. **Pair-Wise Unit Distances:** The distances between every pair of points (in every modality) is also passed
107 into the same modified UMAP distance metric for construction of each modality-specific graph before
108 re-weighting (Fig. 1I). For example, the edge between a pair of units i and j in waveform shape-space is
109 $\theta_{\text{Wave}}(i, j)$. In this way, subsequent calculations use UMAP's notion of distance which has been shown to
110 capture underlying low-dimensional manifold structure (McInnes et al., 2018).
- 111 6. **WNN Construction and Visualization:** Finally, we calculate a connectivity matrix of new pair-wise
112 edges by taking the weighted linear combination of edges in the waveform space and ISI space with the
113 pair-wise distances and affinity weights calculated in the previous steps (Fig. 1J). The final WNN graph
114 is derived from this matrix by using k -nearest neighbors algorithm with k set to a large number (here
115 it is 200). Using UMAP's force-directed graph layout procedure, we then visualize the high-dimensional
116 multi-modal structure by projecting the graph into two dimensions.

117 These steps proceed similarly in the case of three or more modalities with affinity ratios calculated for each
118 modality against every other. The summations in the denominators of Fig. 1H and weighted average of Fig. 1J
119 expand to include these other modalities as well. In the next sections, we use this PhysMAP approach and show
120 how it can combine multiple modalities to delineate cell types in three different datasets.

121 *PhysMAP combines electrophysiological modalities to uncover cell type and laminar structure*

122 *Ex vivo* experiments have shown that the relationship between cell type and electrophysiology spans many different
123 physiological properties (Gouwens et al., 2020; Krimer et al., 2005; Zaitsev et al., 2008). These results are very
124 promising and suggest that cell types could, in principle, be identified purely from modalities available from *in*
125 *vivo* extracellular electrophysiological recording.

126 We examined if combining multiple electrophysiological modalities can find a structure that aligns to cell types
127 better than any modality alone by applying PhysMAP to *in vivo* juxtacellular recordings from the mouse so-
128 matosensory (whisker barrel) cortex (Yu et al., 2019). Juxtacellular recordings provide low noise and well-isolated
129 somatic waveforms recorded unambiguously from single neurons. This dataset also included anatomical and
130 molecular discrete cell type information obtained through measured probe insertions, optogenetic tagging, and
131 immunohistochemistry (Kvitsiani et al., 2013) allowing us to assess if PhysMAP can produce representations
132 that align with ground-truth labels.

133 We first examined the two-dimensional embeddings produced by applying UMAP to these modalities individually
134 (roughly equivalent to applying PhysMAP to single modalities). We found that UMAP applied to waveform
135 shape (the approach taken by WaveMAP; Lee et al. 2021, 2023), led to embeddings that align quite well with
136 underlying cell types and laminar structure (Fig. 2A) as each unit, labeled by cell type and layer, segregates
137 accordingly. However, in this waveform modality, some cell types (such as layer 4 excitatory cells [E-4], in orange)
138 are split in their structure. This split disappears in other modalities such as PSTH (Fig. 2B) or ISI distribution
139 (Fig. 2C). In these other two representations, the structure delineating certain other cell types are lost. For
140 example, PV⁺ and SOM⁺ cells are less clearly delineated from other populations in both the PSTH and ISI
141 distribution visualizations. These results suggest that different modalities provide differing views into underlying
142 structure and perhaps, by optimally combining several modalities, one can find a more robust delineation of cell
143 types.

144 To assess if multimodal combination is useful, we examined the two-dimensional representation of PhysMAP
145 colored according to cell type and layer. We find that a combined representation obtained by weighting each
146 modality on a per-neuron basis appears to best separate cell types (Fig. 2E). The resultant multi-modal repre-
147 sentation also offers additional insight even in the two-dimensional projection: In this visualization, we see that

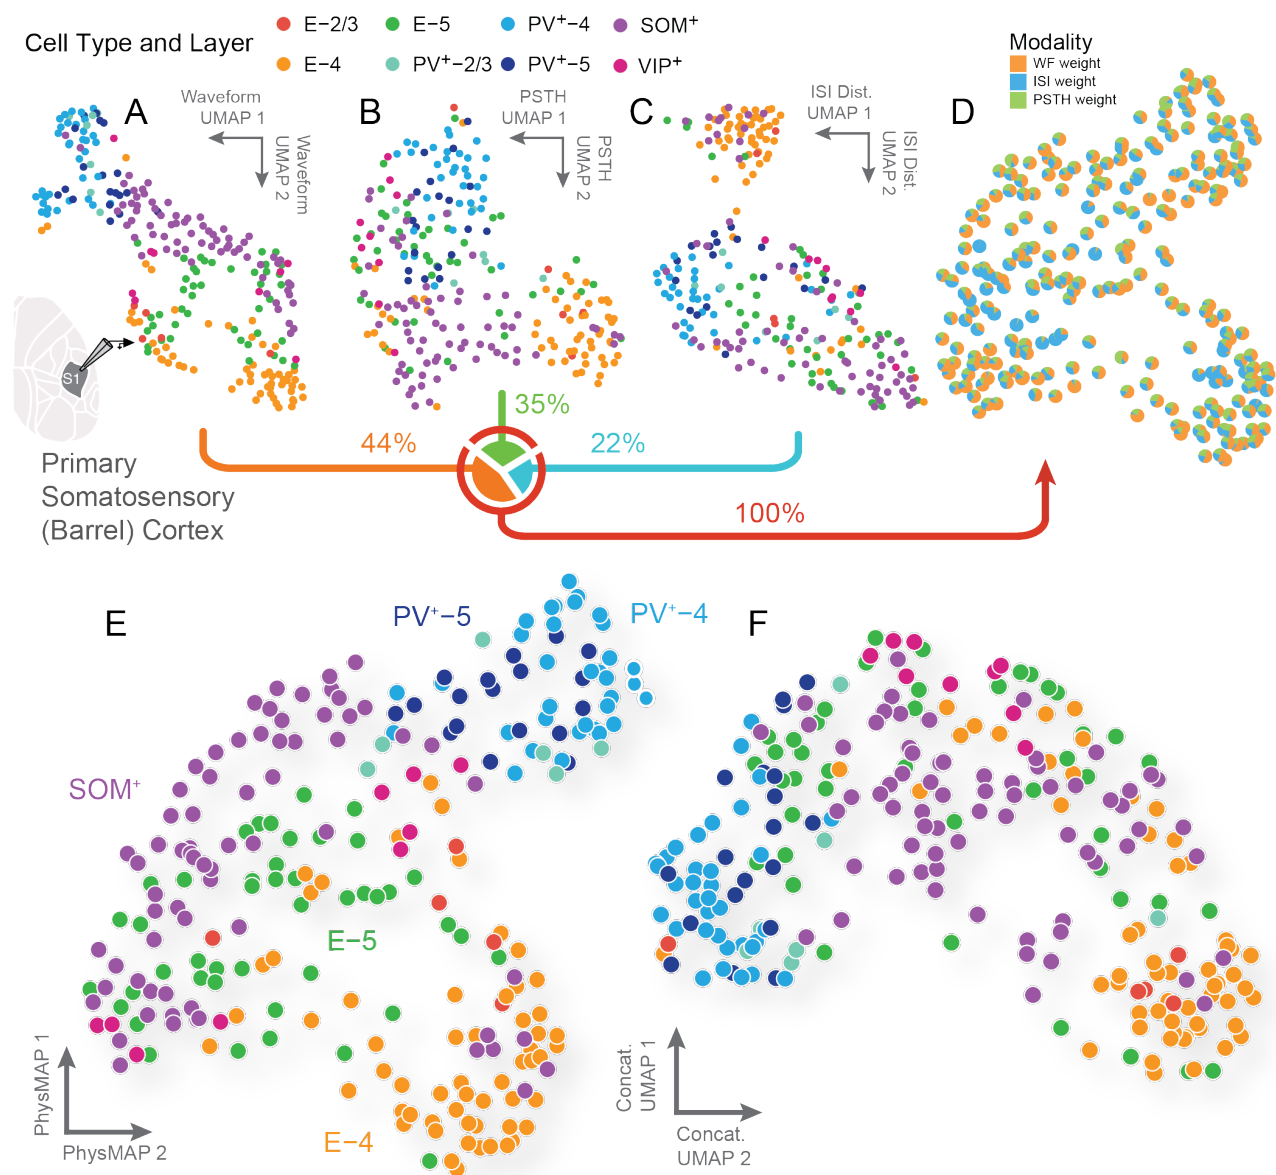


Figure 2: Multi-modal combination of physiological properties leads to structured representations that align with cell types. (A) UMAP on normalized average waveform shapes of neurons recorded *in vivo* juxtacellularly from mouse primary somatosensory cortex (Yu et al., 2019). Each unit is colored according to its ground truth cell type and layer listed above. Each modality was combined in different proportion according to the sum of each modality's per-unit contribution to the total; this total contribution is listed as a percentage below each scatter plot. (B) UMAP on ISI distributions of the same neurons in (A). (C) UMAP on the PSTH firing rate distributions from each neuron in response to a deflection of a whisker during whisking. (D) Each neuron in PhysMAP's WNN representation weighs each modality differently; these modal contributions are the sum of nearest neighbor edge modality weights (β_{Modality} in Fig. 1J) and are shown as proportions of each pie chart. (E) Each of the three modalities are integrated into one representation using the weighted nearest neighbors approach (Hao & Hao et al. 2021) and shown. Again, each neuron is colored according to their ground truth cell type and layer. (F) All three modalities in the dataset from Yu et al. (2019) were concatenated into a single data vector per unit and passed into UMAP. This represents each modality in unweighted combination (as opposed to the WNN approach in PhysMAP). Each neuron in the projected representation is again labeled according to cell type and layer.

148 fast-spiking parvalbumin-positive cells (PV⁺) are tightly clustered and segregated by layer. This visualization
149 also shows that while many SOM⁺ cells are separable they occupy a large area in the PhysMAP representation
150 which implies underlying physiological diversity (Urban-Ciecko and Barth, 2016; Wu et al., 2023; Yavorska and
151 Wehr, 2016). Similarly, we observe that excitatory cells are highly distinct between layers 4 and 5.

152 To gain insight into how the WNN approach works, we next examined how PhysMAP combines each of these
153 modalities on a per-neuron basis. Fig. 2D plots the proportion of each modality used for each data point in
154 Fig. 2A, B, and C after combining all three modalities in PhysMAP and projecting into two dimensions. For
155 the majority of units, waveform shape is the most informative modality as shown by the predominance of its
156 color (orange) across most units. Tallying the contribution of waveform shape, it composes 44% of the total
157 as opposed to PSTH which provides 35% and ISI dist. with 22%. However, there are certain regions of the
158 projection in which other modalities (such as ISI distribution at bottom center of the plot) are more important.
159 Together, these visualizations and results suggest, but do not confirm, that PhysMAP takes a nuanced per-neuron
160 approach to combining modalities and this may help better separate cell types.

161 One concern is that perhaps our PhysMAP representations are a trivial result of having more information through
162 the simultaneous inclusion of multiple modalities. To test if this was the case, we concatenated all three modalities
163 into a single feature vector and ran PhysMAP on this concatenated “super property” which can also be thought
164 of as all three modalities in unweighted combination (Fig. 2F). We found that the representation obtained by
165 simple concatenation of the physiological properties alone demonstrated far weaker structure than the one found
166 in PhysMAP (compare Fig. 2F to Fig. 2E). E-5 and SOM⁺ neurons are more intermixed with one another
167 and also with PV⁺ types. Furthermore, all cell types form less well-structured clusters. This visualization shows
168 that PhysMAP’s improvement in segregating neurons by cell types and layer is not a trivial result due to more
169 information and instead that per-neuron weighted averaging is better in making use of multiple modalities.

170 We emphasize that the two-dimensional representations in this section serve purely as *visualizations* of high-
171 dimensional structure and as such we only examine them here qualitatively. In the following sections, we analyze
172 higher-dimensional representations directly to perform quantitative comparisons.

173 *PhysMAP delineates cell types better than any modality alone or in unweighted combination*

174 The results from the previous section showed that, visually at least, the multi-modal representations obtained
175 from PhysMAP align better with cell types than representations generated by individual modalities alone or
176 in unweighted combination. To more precisely quantify these results in each case (uni-modal, unweighted
177 multi-modal, and weighted multi-modal), we conducted two analyses, one supervised with a classifier and one
178 unsupervised with clustering.

179 In a supervised analysis, we trained a classifier to identify each cell type using either the individual modality
180 representations, the unweighted concatenation, or the weighted PhysMAP representation. For comparison,
181 we also trained classifiers on derived electrophysiological metrics and the full dimensional raw data without
182 dimensionality reduction. See [Methods: Classifier Analysis](#) for more details. For these classification analyses, we
183 projected the WNN or unimodal UMAP graphs into a 10-dimensional embedding space for each of the various
184 modality representations. We conducted an 15-85% test-train split and trained a gradient boosted tree model
185 (GBM) with five-fold cross-validation on this dataset to identify one of five cell types that had over 25 examples
186 each and this process was repeated 10 times with different random seeds. For each of the weighted multi-
187 modal (PhysMAP), unweighted multi-modal (concatenation), and three individual uni-modal representations, we
188 trained this classifier on a 10-dimensional embedding of the WNN graph. We plotted the performance of each
189 of these classifiers trained on the high-dimensional representations against each other. PhysMAP matched or
190 exceeded the performance of all other modalities singularly or in unweighted combination for all cell types (Fig.
191 3A), fully consistent with the visualizations in Fig. 2D and E. We point out that the waveform modality (WF)
192 performed nearly as well as PhysMAP and equating performance for three cell types (PV⁺-4, E-4, and SOM⁺)
193 which confirms our previous observation that waveform shape is the most cell type informative modality (Fig.
194 2D). This finding of waveform shape’s importance in classifying cell types also serves as validation of our previous

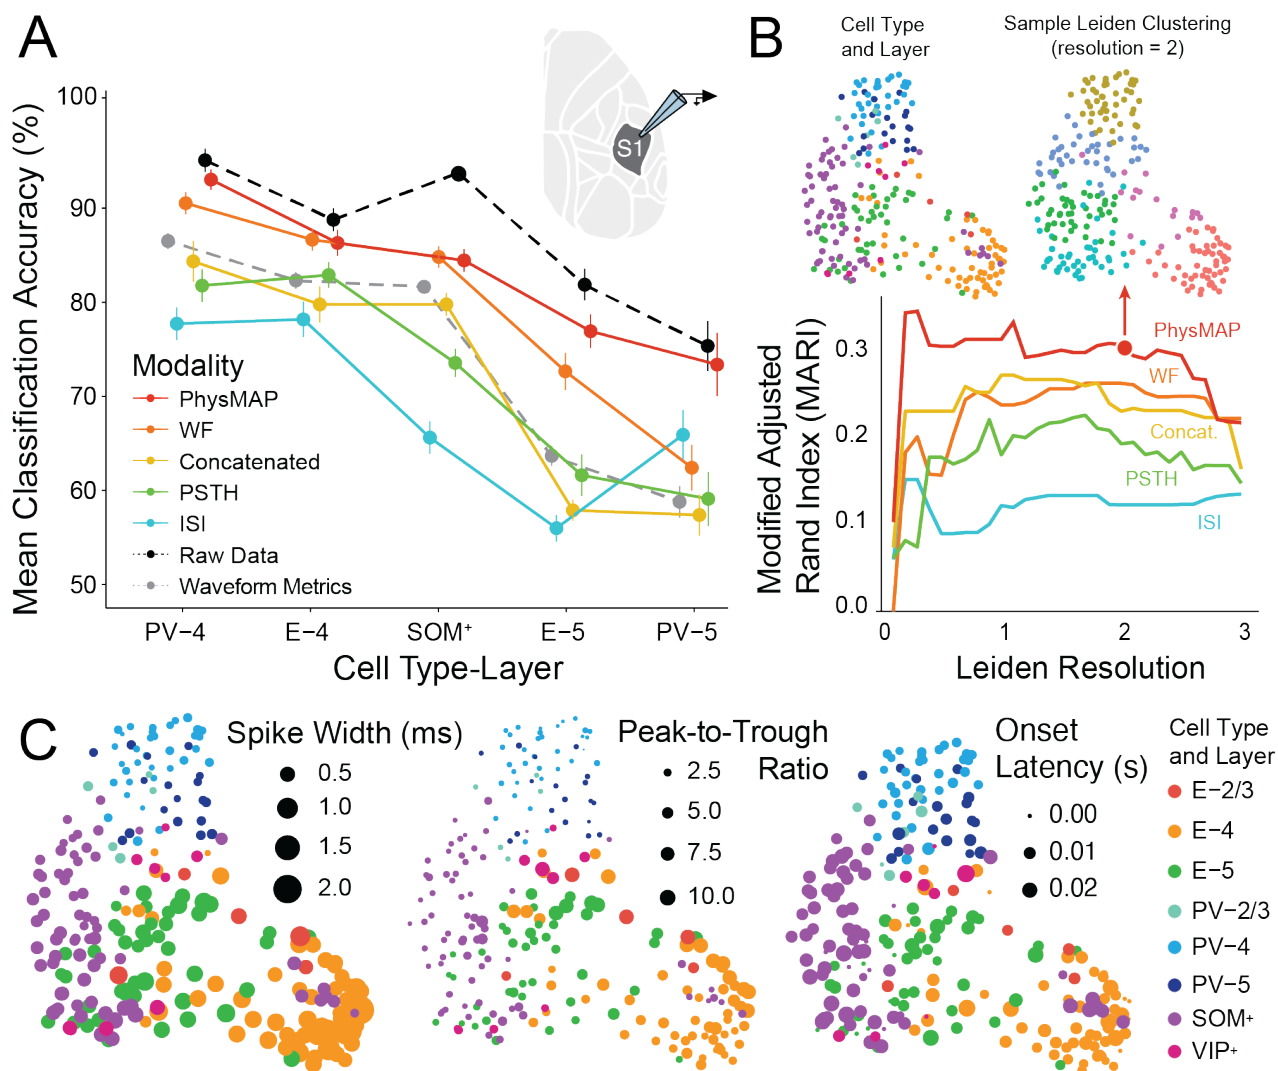


Figure 3: **PhysMAP identifies cell types from juxtacellular recordings better than any modality alone.** (A) A gradient boosted tree model (GBM) classifier was trained with 5-fold cross-validation on the 10-dimensional projection of each modality's UMAP graph individually or on the 10-dim. projection of the multi-modal WNN graph. The same classifier was also trained directly on the full data (that is, without constructing a UMAP graph and projecting it), on the 10-dim. projection of the UMAP graph of the data concatenated into a single feature, and on the two derived waveform metrics (spike width and peak-to-trough ratio). The balanced accuracy performance of this classifier (mean \pm S.E.M.) on held-out data for each modality, combined modalities, and derived metrics is shown for the five cell type classes with over 25 units each. (B, top) The PhysMAP projection of neurons with their ground truth cell type and layer (left) next to an example Leiden clustering (Traag et al., 2019) with resolution parameter set to 2. (B, bottom) Leiden clustering is applied to the UMAP graphs of each modality alone (waveform [WF], ISI dist., and PSTH), in weighted (PhysMAP), or unweighted combination (concat.) and shown with the associated modified adjusted Rand index (MARI; Sundqvist et al. (2020)) calculated across a range of resolution parameter values from 0.1 to 3.0 in 0.1 step increments. Arrow marker indicates the clustering on PhysMAP used above in B. Waveform metrics and raw data were omitted in this analysis because these do not yield a high-dimensional UMAP graph. (C) The spike width, peak-to-trough ratio, and spiking onset latency for each neuron are shown via log transformed marker size under PhysMAP's two-dimensional projection.

195 approach and findings with WaveMAP (Lee et al., 2021). The benefit of adding the additional modalities comes
 196 from the improved classifiability of the E-5 and PV⁺-5 neurons from the rest, which was much poorer with only
 197 waveforms.

198 PhysMAP also demonstrated better classification performance than the modalities in an unweighted combination
199 (concatenated) for all cell types. Similar to visualizations comparing the embeddings of the weighted and un-
200 weighted modalities in combination (Fig. 2E vs Fig. 2F), our classification analysis demonstrates that PhysMAP
201 is able to draw out cell type structure in a manner that is not trivially the result of including several modalities
202 simultaneously.

203 Finally, PhysMAP (and the full waveform) far outperform derived waveform metrics which are how cell types are
204 traditionally identified. This is a result that is consistent with general notions in machine learning that hand-
205 derived metrics, while seemingly “interpretable” can lose important information. In our case features do not fully
206 capture the full diversity of waveform shape in a way that is useful for identifying cell types.

207 **Control:** Classification performance from PhysMAP was robust to both the embedding dimensionality chosen for
208 the WNN graph. First, based on a linear estimate of intrinsic dimensionality using principal component analysis, a
209 10-dimensional embedding was chosen for the classification analysis given that the “elbow”, normally regarded as
210 the sufficient number of dimensions to capture dataset variance, was at the third, fourth, and third PC’s (which
211 sum to 10) representing 91%, 80%, and 93% of the total variance for waveform shape, ISI dist., and PSTH
212 respectively (Fig. S1). However, we also found that the particular choice of embedding dimensionality actually
213 did not have a large effect on classifier performance. We examined the full range of embedding dimensionalities
214 from 30 (matching the lowest ambient dimensionality of the input modalities) down to 2 (the plot Fig. 2E itself)
215 and found very similar patterns of classification accuracy across all cell types and all modalities (Fig. S2A).

216 **Control:** Similarly, the particular classification algorithm used had little effect on this classification analysis. We
217 also performed this supervised classification analysis across five other classifiers: a random forest, radial basis
218 support vector machine (rbSVM), a classification tree, a k -neighbor nearest classifier, and a neural network with
219 single hidden layer. Each showed very similar results demonstrating that our results are not classifier-dependent
220 (Fig. S2B). Thus, in Fig. 3A we demonstrate that PhysMAP is able to intelligently combine multiple modalities
221 in a manner that draws better cell type delineations than any modality alone or in unweighted combination.

222 **Comparison:** One potential criticism is that the interpretability of our PhysMAP approach comes with a cost
223 and that some information is lost. However, PhysMAP performs nearly as well as a classifier on the raw data
224 in detecting cell types. When the GBM classifier is trained directly on the full dimensional dataset from all
225 modalities directly (that is, not on a low-dim. projection of the UMAP graph), performance is not much better
226 on most cell types except for SOM⁺ cells (Fig. 2A, dashed line). Furthermore, training a classifier directly on
227 the data loses all interpretability since it only returns a probability that some neuron belongs to a given class.
228 Thus, while there is some information loss because of PhysMAP’s dimensionality reduction, it still produces
229 a representation that contains most of the information necessary for both cell type classification. In return it
230 provides a useful visualization for exploring physiological diversity. We will explore how this is useful in the next
231 section.

232 *PhysMAP is also better in unsupervised settings*

233 We also examined if PhysMAP representations aligns to underlying cell types even when no labels are available
234 to train on. In most electrophysiological datasets, only a subset of the recorded cells will have ground truth
235 labels. In a fully-unsupervised analysis, we asked the question of “without access to *any* ground truth, how well
236 do unsupervised methods identify cell types?”. Thus, by examining the representations produced by PhysMAP in
237 the extrema of having both complete ground truth or zero ground truth (supervised vs. unsupervised algorithms),
238 we gain a more complete picture of PhysMAP’s performance over a range of possible dataset conditions.

239 We assessed how well unsupervised clusters aligned to underlying ground truth cell types by calculating the
240 modified adjusted Rand index (MARI; Sundqvist et al. (2020)) between each. This is a measure of how often the
241 same ground truth cell types co-cluster under an unsupervised clustering. We used cluster labels generated from
242 Leiden clustering applied to the high-dimensional graphs produced by PhysMAP and each physiological modality
243 (that is, the graphs before projection) across a range of resolutions (cluster sizes) and calculated the MARI

244 score between the two sets of labels. See [Methods: MARI Calculation](#) for more details. The greater the MARI
245 value, the more the clusters contain singular cell types each. To provide an example of how Leiden clustering on
246 PhysMAP compares to ground truth, we show neurons in PhysMAP colored according to cell type ([Fig. 3B, top](#)
247 [left](#)) alongside a sample Leiden clustering with resolution set to 2 ([Fig. 3B, top right](#)). Examining each resolution
248 parameter of Leiden clustering, PhysMAP produced higher MARI scores than any electrophysiological modality
249 alone or in unweighted combination for all Leiden resolutions ([Fig. 3B](#)). This performance was followed by
250 modalities in unweighted combination and waveform shape alone, PSTH, and then ISI distribution. Collectively,
251 these results demonstrate that PhysMAP is useful in identifying cell types in both supervised and unsupervised
252 settings.

253 *PhysMAP provides interpretable representations and biological insight*

254 Thus far, we have demonstrated that PhysMAP outperforms individual modalities at uncovering the link between
255 electrophysiological properties and underlying cell types in supervised and unsupervised settings. However, it is
256 increasingly clear that cell types are not “discrete” and physiological variation within a cell class is the norm and
257 not the exception ([Cembrowski and Menon, 2018](#)). Ideally, any tool for cell type identification should allow one
258 to visualize this diversity, both within and across discrete cell types.

259 We found that PhysMAP’s visualizations were ideal for this purpose of examining underlying diversity. [Fig. 3C](#),
260 shows the PhysMAP plot in [Fig. 2E](#) but now with marker radius scaled to either reflect waveform spike width,
261 peak-to-trough ratio, and onset latency. Examining spike width [Fig. 3C, left](#), many trends immediately pop out:
262 PV⁺ cells seem to be uniformly narrow no matter the layer; SOM⁺ cells occupy a gradient of widths and overlap
263 with E-5 cells at their widest; and the widest spiking cells are E-4. This finding regarding SOM⁺ cells occupying
264 a range of spike widths is likely because it is composed of several subtypes of varying widths ([Hostetler et al.,](#)
265 [2023](#)) which have distinct computational roles ([Wu et al., 2023](#)). Waveform peak-to-trough ratios also differ
266 between cell types [Fig. 3C, middle](#) with PV⁺ and SOM⁺ cells having the smallest and both E-4 and E-5 cells
267 occupying a range which includes the largest ratios. Finally, we replicated that onset latency (half the time it
268 takes to reach peak firing rate after whisker deflection; [Fig. 3C, right](#)) is fastest for PV⁺ cells and SOM⁺ cells
269 are some of the slowest, consistent with the reports from the original study ([Yu et al., 2019](#)).

270 These insights were less clear than when only examining the metrics without using the PhysMAP representation.
271 We examined a scatter plot of the spike width and peak-to-trough-ratio to assess whether any structure easily
272 pops out ([Fig. S3](#)). While the scatter plot shows some cell type structure it appears like one continuous
273 relationship and fails to show any easy separability between SOM⁺, PV⁺ and excitatory cells. In addition, in
274 either waveform shape metric, the distributions of the the PV⁺ and excitatory cells of different layers overlap as
275 shown by the coinciding PV⁺-4/PV⁺-5 or E-4/E-5 histograms on the marginals of [Fig. S3](#).

276 In summary, PhysMAP provides the ability to easily and simply combine multimodal electrophysiological data
277 to obtain a representation that 1) classifies cell types better than any individual modality or feature alone, 2)
278 can also be used in unsupervised settings to separate out candidate cell types, and 3) provides an interpretable
279 visualization that can be used for additional exploration and development of hypotheses about data.

280 *PhysMAP uncovers cell types in extracellular electrophysiological recordings*

281 The juxtacellular recordings analyzed in the previous section represent a best case scenario for PhysMAP for several
282 reasons. First, this recording method offers access to somatic waveforms that are unperturbed by distance to
283 spike origination site or from differences in spike shape occurring due to the differential ion channel properties
284 of neurites (axons and dendrites) ([Someck et al., 2023](#)). Second, juxtacellular recordings offer unambiguous
285 isolation of single unit activity. In contrast, extracellular recordings are prone to poor isolation which blurs the
286 properties of different cell types ([Vincent and Economo, 2023](#)). This results in many multi-unit “cell types”
287 occupying intermediary properties. Third, juxtacellular recordings allow for the preferential targeting of particular
288 cell types that may be less prevalent. Inhibitory types make up only 20% of neurons in mouse cortex ([Tremblay](#)
289 [et al., 2016](#)) but in this set of juxtacellular recordings they compose 56% of the dataset as they have been

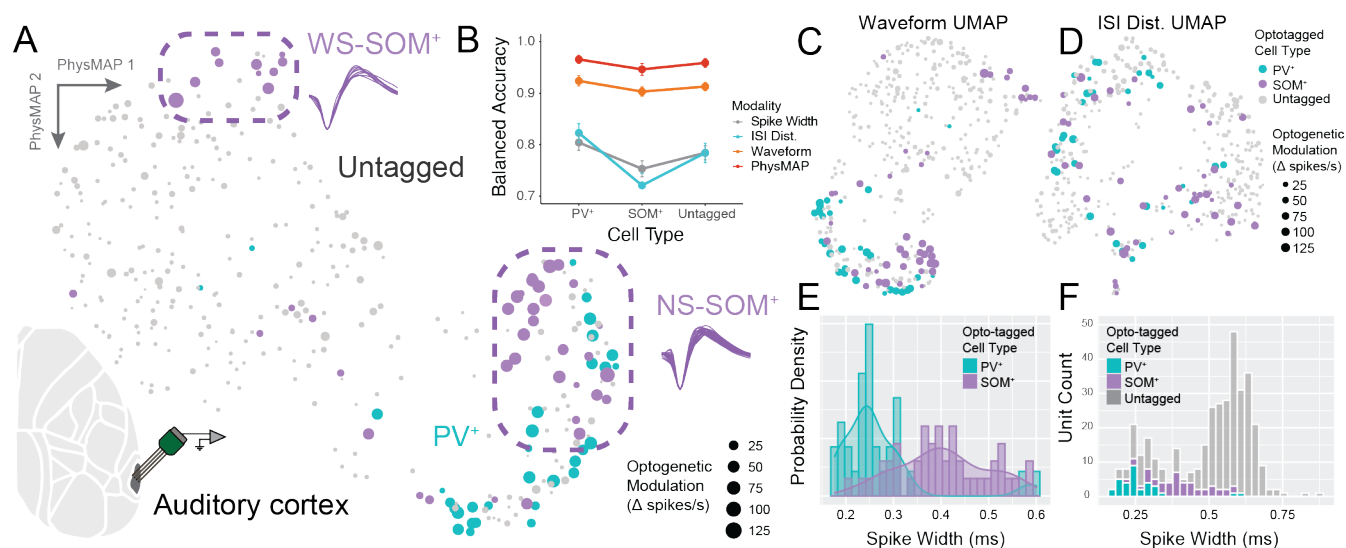


Figure 4: PhysMAP identifies cell types in extracellular recordings. (A) PhysMAP applied to extracellular recordings from mouse auditory cortex (Lakunina et al., 2020) using silicon probes. We used waveform shape and ISI distribution as our modalities and neurons are colored according to cell type whose identities were obtained via optogenetic tagging. Marker size is set according to the increase in spike rate from optogenetic stimulation. Highlighted are a population of wide-spiking SOM⁺ cells (top) and narrow-spiking SOM⁺ cells (right). (B) A GBM classifier with five-fold cross-validation was trained on the 10-dimensional graph of each modality, PhysMAP, or waveform spike width to identify each optotagged cell type or untagged label. The balanced classification accuracy is shown (mean \pm S.E.M.) with many error bars smaller than marker size. (C) Normalized average waveforms in this dataset are passed into UMAP and their projection shown with optotagged cells colored and degree of optogenetic modulation (average change in spikes/s during stimulation epochs) shown by marker size. (D) Similarly, ISI distributions for these same neurons are also passed into UMAP and their projection shown. (E) Probability density estimates for the distribution of each optotagged type across a range of spike widths is shown both as a histogram and with a kernel density estimator (solid line). (F) Additionally, spike width is shown for all cells, including untagged ones.

290 preferentially targeted for recordings. This provides the downstream benefit in that a classifier will not be biased
 291 towards more prevalent cell types or will fail to classify less common ones. We asked whether PhysMAP could
 292 identify find cell types even when applied to recordings acquired via extracellular probes. To address this question,
 293 we first analyzed an extracellular electrophysiological dataset collected from mouse primary auditory cortex (A1,
 294 Lakunina et al., 2020) with optotagged cells.

295 In this experiment, Lakunina et al. (2020) recorded from neurons in A1 with extracellular probes while mice
 296 were presented with auditory stimuli. They identified both SOM⁺ and PV⁺ units using optogenetic tagging and
 297 recorded each unit's waveform shape and ISI distribution (over 50 ms windows). We passed these two modalities
 298 into PhysMAP to produce a combined representation (Fig. 4A). In this projected space, both SOM⁺ and PV⁺
 299 cell types (purple and teal respectively) are well-separated and occupy a distinct region of the embedding that
 300 clusters away from a much larger number of "untagged" (presumably excitatory) cells (in gray, Fig. 4A). This
 301 result complements the findings in the previous juxtacellular dataset (Yu et al., 2019) and confirms that PhysMAP
 302 is able to separate out cell types even in an extracellular setting.

303 We again examined if PhysMAP better separates out these cell types than any single modality or waveform
 304 metrics alone. As in the previous section, we quantified this separation by training a classifier to identify each
 305 cell type using PhysMAP and benchmarked this classifier's performance against classifiers trained on individual
 306 modalities (waveform shape or ISI distribution) or waveform metrics (spike width). Note that MARI was not
 307 calculated for this dataset as such an analysis is only valid when all labels are known. The classifier on PhysMAP
 308 identified PV⁺ and SOM⁺ cells with high accuracy (97% and 95% mean balanced accuracy, respectively) and

309 surpassed classification performance using ISI distributions or waveform metrics (82% and 72% or 80% and
310 75% mean balanced accuracy, respectively), the latter of which are how cell types are identified in much of the
311 literature (Fig. 4B). PhysMAP also slightly outperformed classification on waveform shape for both PV⁺ and
312 SOM⁺ cells (92% and 90% mean balanced accuracy, respectively). The reason for this near-equivalence is that
313 waveform shape is highly informative of underlying cell types and thus PhysMAP correctly prioritizes this modality.
314 The median unit weight associated with the waveform shape and ISI distribution modalities were 0.71 and 0.29
315 respectively. This prioritization of waveform shape over ISI distribution was also evident in their respective UMAP
316 visualizations: the PhysMAP's projection appears almost exactly the same as UMAP on waveform shape (Fig.
317 4C). In contrast, a UMAP applied to ISI distribution shows far less structure Fig. 4D. Furthermore, although
318 these does not correspond to a particular cell type, classification with PhysMAP outperforms any other modality
319 in identifying “untagged” cells. By mirroring results from our previous analysis of juxtacellular recordings, we
320 establish that PhysMAP is also capable of uncovering cell types in extracellular recording settings.

321 Again, as in the juxtacellular dataset, the PhysMAP approach provided interpretable representations. To illustrate
322 this point, we identified two sub-populations of SOM⁺ cells not documented in the original publication of this
323 dataset (Lakunina et al., 2020) but predicted by the literature (Hostetler et al., 2023; Wu et al., 2023; Urban-
324 Ciecko and Barth, 2016). These wide-spiking SOM⁺ (WS-SOM⁺, Fig. 4A, upper) and narrow-spiking SOM⁺
325 (NS-SOM⁺, Fig. 4A, lower right) sub-populations occupied separate regions of the PhysMAP embedding and
326 differed significantly in their spike width (0.48 ± 0.04 vs. 0.37 ± 0.05 ms [median \pm S.D.], respectively; p
327 < 0.001 , Mann-Whitney U test). These sub-populations have been observed in other experiments to pertain
328 to SOM⁺ cellular subtypes with specific synaptic targeting (Wu et al., 2023) and differential function during
329 behavior (Kvitsiani et al., 2013; Kim et al., 2016).

330 Note that these two sub-populations are again not trivially identifiable from spike width alone: SOM⁺ cells, as a
331 whole, were part of a continuum under this metric (Fig. 4E, in purple) and overlapped with PV⁺ cells (Fig. 4E,
332 in teal) in width. Even if an arbitrary width cutoff was used to bisect SOM⁺ cells into two sub-populations, the
333 WS-SOM⁺ cells would occupy a range of widths that strongly coincides with large number of other untagged—
334 presumably excitatory—cell types (Fig. 4F) and thus would not be able to be identified uniquely. Only PhysMAP
335 is able to simultaneously disambiguate these two SOM⁺ subtypes in addition to PV⁺ and untagged putatively
336 excitatory cells.

337 *PhysMAP finds the same cell types across heterogeneous datasets from different labs*

338 We have demonstrated that PhysMAP uncovers cell types in datasets from single labs but it would be beneficial if
339 a classifier could also find cell type structure across datasets gathered by different labs. Cell type classifiers could
340 then be built by small collaborations of interested investigators across different labs. To this end, we combined
341 extracellular recording datasets originating from separate labs each containing cells from both visual cortex and
342 hippocampus (Siegle et al., 2021; Petersen et al., 2020) released as part of CellExplorer (Petersen et al., 2021).

343 For neurons from both datasets, we had measurements of waveform shape, ISI distribution, spike autocorrelogram
344 (ACG), and eleven derived electrophysiological metrics such as spike width, coefficient of variation, and ACG
345 delay time constant etc (full list in [Methods: Extracellular Mouse Visual Cortex and Hippocampus Dataset](#)).
346 Similar to our two previous analyses, all modalities ostensibly showed cell type-dependent structure but differed
347 in which cell types were better separated: waveform shape (Fig. 5A) splits PV⁺ cells into two regions whereas
348 the other modalities do not. Similarly, pyramidal cells are well-isolated under ISI distribution (Fig. 5B) but less
349 well in others like ACG (Fig. 5C). Derived electrophysiological metrics (Fig. 5D) present an intuitive PV⁺ to
350 SOM⁺ to pyramidal cell structure but separate out SOM⁺ cells less than using the ISI distribution. PhysMAP
351 combines these modalities to produce a representation that produces a clear separation of cell types (Fig. 5E).

352 In order to evaluate PhysMAP quantitatively for this dataset, we again calculated both the accuracy of a classifier
353 and each representation's MARI score. We again trained a GBM classifier with 5-fold cross-validation to identify
354 the four cell types with more than 25 units each from the 10-dimensional embedding of each modality and

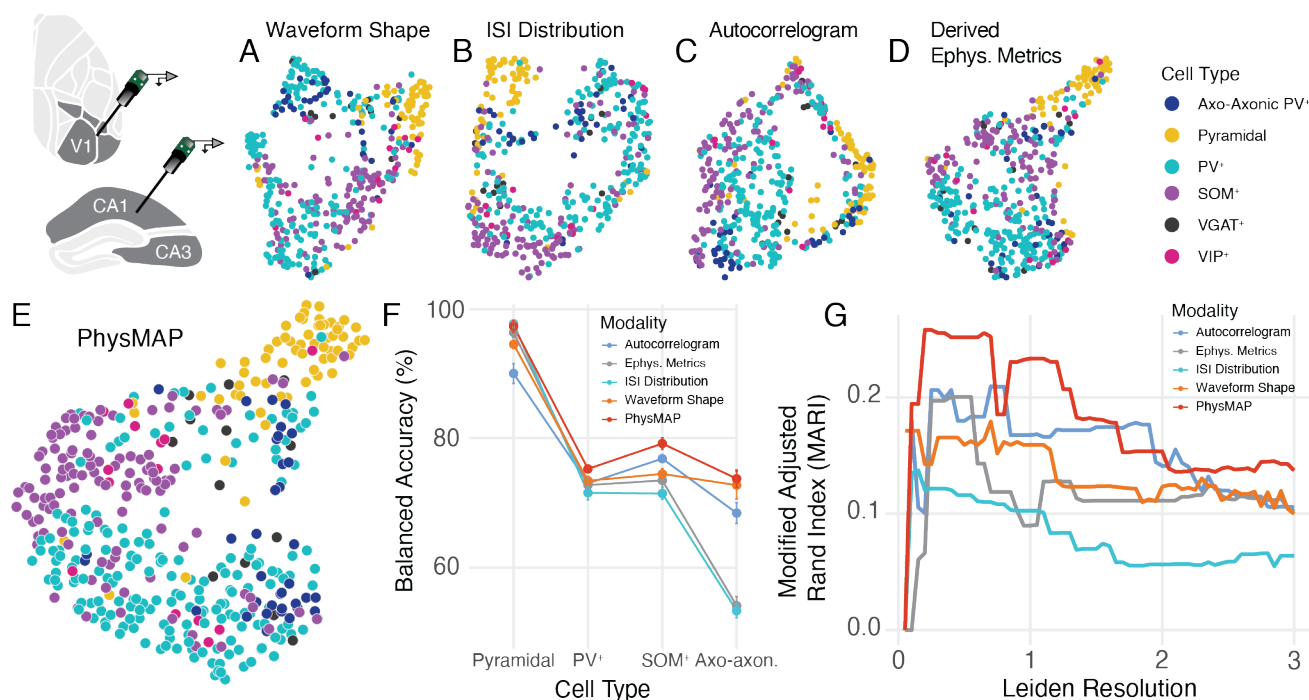


Figure 5: **PhysMAP is able to identify cell types across areas and labs.** (A - D) Across a pooled dataset of visual cortex and hippocampal extracellular recordings from CellExplorer, we show 2-D UMAP visualizations applied to waveform shape, inter-spike interval distribution, autocorrelogram, and thirteen various derived electrophysiological metrics respectively. Cell type classifications were obtained by opto-genetic tagging. (E) The previous four modalities are combined into a unified PhysMAP representation and visualized in 2D. (F) A gradient boosted tree model (GBM) classifier with 5-fold cross-validation was trained on the 10-dimensional projection of each modality's high-dimensional UMAP graph or on PhysMAP's multi-modal WNN graph. The balanced accuracy for the classifier is shown for each modality and cell type that contained more than 25 examples. (G) Leiden clustering is applied to each modality alone and to PhysMAP and the corresponding modified adjusted Rand indices are calculated across a range of resolution parameter values from 0.1 to 3.0 in 0.1 step increments. All cell type labels were included in this calculation.

355 PhysMAP. For all cell types (pyramidal, PV⁺, SOM⁺, and axo-axonic [a PV⁺ subtype]), PhysMAP provided the
 356 best classification accuracy (Fig. 5F).

357 Next we evaluated how well cell types clustered by calculating the MARI score of each representation. As before,
 358 across nearly all Leiden resolution values, PhysMAP provided the best unsupervised clustering of cell types (Fig.
 359 5G). We also examined if hippocampal and cortical inhibitory cells were differentiable (Fig. S4A). Both PV⁺
 360 (Fig. S4B) and SOM⁺ (Fig. S4C) cells did not form separate populations in PhysMAP regardless of the area
 361 despite being well-differentiated between types. Thus, we expect that that a cell type classifier will generalize
 362 across areas to some extent even between areas as different as hippocampus and neocortex. Thus, even when
 363 combining data from multiple labs and areas, PhysMAP better separates cell types than any single modality
 364 suggesting the feasibility of building cell type classifiers in a collaborative manner. We next evaluate how well
 365 PhysMAP performs as a cell type classifier.

366 *PhysMAP allows for the creation of a cell type classifier to accurately identify several major cell types*

367 So far, we have shown that PhysMAP is better able to unsupervisedly identify cell types in three areas of
 368 mouse cortex and hippocampus but in these analyses, we compared modalities by training a classifier after data
 369 transformation to facilitate comparison with limited data sizes. In order to properly test PhysMAP's functioning
 370 as a cell type classifier, we evaluate classification on held-out, un-transformed data. To build the cell type
 371 classifier, we took advantage of Seurat's multi-modal reference mapping tool which allows for classification of

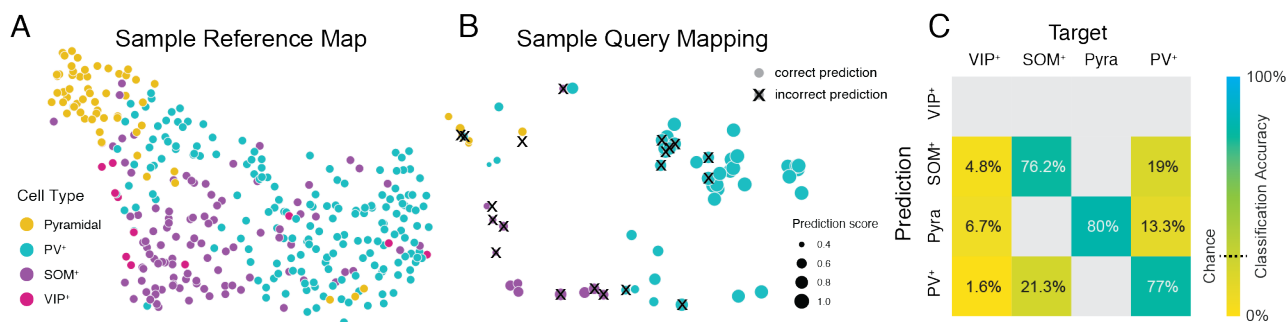


Figure 6: **PhysMAP makes it possible to construct a cell type classifier for three major cell types.** (A) Using the previous CellExplorer dataset, axo-axonic cells were relabeled as PV⁺ (because they are a PV⁺-subtype) and VGAT⁺ cells were omitted since they label inhibitory neurons as a whole. The four modalities in (Fig. 5A-D) were then passed into PhysMAP for a given training set and a 2-D visualization for this “reference map” is shown. (B) The held-out test set is then projected into the reference map space using anchor alignment and visualized as a “query mapping”; cell type labels are according to nearest neighbor predictions after projection. Incorrect predictions identified by a small “x” plotted over each data point. Each cell is also sized according to “prediction score” which is an empirical measure of the quality of nearby anchors; greater prediction scores imply a closer match between local query mapping and the reference map at that location. (C) A confusion matrix of query mapping predictions used as a classifier across the four major cell types in (A). Each row contains the mean percentage (after 10 instances of classifier training) that each predicted cell type label was assigned to its true target cell type. The percentage of cells that were correctly predicted is shown along the main diagonal. Some tiles are left blank because no predictions were made for that pair of predicted and target cell type.

372 cell types given some reference dataset. Here, we take the previous CellExplorer dataset and test PhysMAP’s
 373 ability to identify the identity of a random subset of the total dataset with the remainder used to build a “reference
 374 map”. With this analysis, we show that PhysMAP maintains good generalization on three of four major cell types
 375 available (excitatory pyramidal, PV⁺ interneurons, and SOM⁺ interneurons).

376 To begin, we consolidated the CellExplorer dataset such that axo-axonic PV⁺ chandelier cells were relabeled
 377 as “PV⁺” (Dudok et al., 2021) and removed VGAT⁺ neurons as they are pan-inhibitory and thus do not label
 378 one particular inhibitory cell type. We then divided the 417 resulting cells into an 15-85% test-train split. The
 379 training data—containing waveform, ISI distribution, and single derived metrics measuring the former two—was
 380 passed into PhysMAP to generate a reference map (Fig. 6A). A query dataset of held-out neurons was then
 381 mapped onto this reference using anchor point alignment (Haghverdi et al., 2018) and cross-projection of query
 382 data onto the reference map (Stuart & Butler et al. 2019). See [Methods: Reference Mapping with Anchors and](#)
 383 [Cell Type Classification](#) for methodological details.

384 Fig. 6B shows a sample query mapping of the test set colored according to predicted cell type with incorrect
 385 predictions overlaid with a black “x”. Cell type predictions are made based upon a query cell’s nearest neighbors
 386 in the reference map. These data points are also sized according to their “prediction score” which is a per-neuron
 387 measure of confidence in the anchor alignment; higher scores are associated with higher quality alignment. In
 388 this sample query mapping, neurons with low prediction score more often result in incorrect predictions. Thus,
 389 if a more conservative prediction of cell types is desired, only neurons with higher prediction scores can be used.

390 To evaluate the anticipated performance on this dataset with PhysMAP performing as a cell type classifier, we
 391 conducted ten random test-train splits of the dataset and examined the average classification performance. We
 392 found that although VIP⁺ cells were not identifiable (the query mapping never predicted them due to low sample
 393 number) all other cell types were highly identifiable (Fig. 6C). Each of SOM⁺, PV⁺, and excitatory pyramidal
 394 cells were identified at or above 75% accuracy which was far above chance (25%). In summary, we find PhysMAP
 395 not only locates more cell type structure than other modalities, it also provides a cell type classifier that can be
 396 used to identify multiple cell types from *in vivo* extracellular recordings.

397 Discussion

398 Our goal in this study was to assess if a method to compute a weighted combination of electrophysiological
399 modalities could identify underlying neuronal cell types and serve as a classifier. Our approach, PhysMAP,
400 integrates electrophysiological modalities to produce a combined representation that separates out cell types
401 better than any single modality alone or in unweighted combination. Specifically, in both juxtacellular and
402 extracellular *in vivo* experiments, the multi-modal representation produced by PhysMAP was better at uncovering
403 cell types in both unsupervised and supervised settings. We also demonstrated excellent classification performance
404 in identifying three major cell types from held-out data in an interpretable manner.

405 We note that PhysMAP is not the only approach for merging information from multiple modalities. Other
406 prominent tools in multiomics including MOFA ([Argelaguet et al., 2020](#)) and MOJITOO ([Cheng et al., 2022](#))
407 can perform similar integration of multiple modalities. We used the WNN algorithm from Seurat v4 to develop
408 our PhysMAP approach because it has demonstrated excellent performance for multiomics problems ([Cheng
409 et al., 2022](#)) and we found that this performance extends to identifying cell types when combining multiple
410 electrophysiological modalities. We also used Seurat because it is perhaps the most extensively used package in the
411 multiomics community, has a large community of contributors, and is continually being improved. For instance,
412 Seurat v5 adds “bridge integration” allowing the prediction of a cell’s proteomic profile given its transcriptomic
413 expression by means of a “dictionary” ([Hao et al., 2023](#)). Applied to an electrophysiological recording, one might
414 envision predicting a neuron’s waveform shape given only knowledge of its ISI distribution.

415 We used the waveform, inter-spike interval, and peri-stimulus time histogram in PhysMAP to differentiate cell
416 types. Other modalities might also be important for delineating cell types. For example, the spike-local field
417 potential (spike-LFP) coupling has been posited to separate PV⁺ and SOM⁺ cell types ([Wei et al., 2023](#); [Onorato
418 et al., 2023](#)). In addition, there may be particular “localizer” stimuli that can be used to generate PSTHs to
419 discriminate certain cell types. For example, a full-field flash has long been used to identify thalamocortical input
420 cells into V1 ([Heynen and Bear, 2001](#)). However, one must be judicious about the modalities chosen because of
421 circularity. That is, if a certain cellular property wants to be compared between cell types, it should not be passed
422 into PhysMAP because such properties will already be used to select the groups to be compared ([Kriegeskorte
423 et al., 2009](#)).

424 PhysMAP offers two advantages over optical imaging approaches: simultaneous monitoring of neuronal cell types
425 on single trials and the exploration of cell types in primates. First, they allow for the *simultaneous* monitoring
426 of neural populations, which is currently difficult in optical imaging experiments due to constraints on optics
427 that limit imaging to only two fluorophores ([Aharoni and Hoogland, 2019](#)). Simultaneous monitoring will help
428 better understand circuit dynamics on single trials ([Peixoto et al., 2021](#); [Boucher et al., 2023](#)), and bridge the
429 gap between anatomical microcircuits and dynamics ([Esparza et al., 2023](#)). Second, cell type-specific imaging
430 often requires transgene delivery and expression. In contrast, PhysMAP only requires recordings of the electrical
431 activity of single neurons and is therefore particularly useful for the study of cell type-specific neural dynamics in
432 non-human primates (NHPs). NHPs are a valuable and scarce resource and each fully-trained animal can perform
433 multiple sophisticated tasks. Primate researchers are therefore cautious in using direct cell type-specific recording
434 techniques that entail risky procedures such as viral injections or cranial window installations. In addition, in
435 primates, optical methods and genetic engineering tools are difficult to execute without enormous resources and
436 the geometric constraints are even more complicated because of the size of cortex ([Trautmann, O’Shea, & Sun
437 et al. 2021, O’Shea et al., 2017](#)). For this reason, PhysMAP would be a no risk solution to obtaining cell types
438 when applied to routine electrophysiology experiments.

439 Naively, one might be tempted to question why dimensionality reduction methods like PhysMAP should be used
440 at all. Instead, would a purely classification-based approach not suffice to identify cell types? While current
441 classifiers will delineate some cell types from electrophysiological properties ([Schneider & Azabou et al. 2023,
442 Ye & Shelton et al. 2023](#)), we find that these do not outperform PhysMAP in identifying the same cell types.
443 More importantly, PhysMAP has two key advantages over **pure discrete classification** approaches.

444 First, PhysMAP produces interpretable visualizations that can reveal heterogeneity within and across cell types.
445 This heterogeneity manifests as an overlap in the embeddings of one cell type over another in the PhysMAP
446 visualization. For example, we observe considerable variability within SOM⁺ cell types in both juxtacellular and
447 *in vivo* extracellular settings and across cortical areas. In both the mouse S1 and A1 datasets, we found that
448 SOM⁺ cells were heterogeneous being composed of both broad- and narrow-spiking waveforms, with some of the
449 SOM⁺ cells co-localizing with excitatory neurons in the PhysMAP representations. This is entirely expected given
450 the known heterogeneity of SOM⁺ cells which have variations in morphology as well as physiological properties
451 relating to underlying sub-cellular types (Wu et al., 2023). In particular, many SOM⁺ cells demonstrate a fast-
452 spiking phenotype in rodents and are known to express fast-spiking potassium channels (Urban-Ciecko and Barth,
453 2016). In contrast other SOM⁺ cells, such as Calb2, exhibit broad-spiking phenotypes (Hostetler et al., 2023).
454 Similarly, in the juxtacellular dataset, excitatory neurons formed a broad class and layer 4 excitatory neurons
455 had different and distinct electrophysiological properties from layer 5 excitatory neurons. While not directly
456 investigated, this is in agreement with findings from intracellular recordings in motor cortex that corticothalamic-
457 (deep layer 5 and layer 6) and intratelencephalic-projecting neurons (layer 2/3 and superficial layer 5) have distinct
458 electrophysiological properties and morphologies (Scala et al., 2021). We anticipate this is at least in part why
459 excitatory cells of different layers segregate in PhysMAP's visualizations. These heterogeneities in SOM⁺ and
460 excitatory cells, among others, are likely because cell types are not discrete and increasingly studies suggest they
461 lie on a continuum. Said differently, variation within a discrete "cell type" is the norm and not the exception
462 and can emerge from differences in morphology, laminar location, and gene expression (Cembrowski and Menon,
463 2018; Scala et al., 2021). This true underlying physiological variation is invisible to typical discrete classifier
464 approaches but readily apparent with PhysMAP.

465 Second, PhysMAP offers a head-to-head comparison of each modality that many purely classifier methods do
466 not provide easily. For instance, in both juxtacellular and extracellular recordings, waveform shape is the most
467 reliable modality for delineating ground truth cell type. These results reaffirm our previous work (Lee et al.,
468 2021) and with the observation that the expression of many genes which determine waveform shape often covary
469 with neuronal cell types in different layers of cortex (Bomkamp et al., 2019). For instance, analysis of a mouse
470 Patch-seq dataset revealed that the gene expression of many different potassium ion channels (e.g., *Kcnn7*
471 and *Kcnc2*) correlated with properties of the waveform such as amplitude of after-hyperpolarization between
472 excitatory and inhibitory cell types (e.g., *Lamp5*, *Pvalb*, etc, Bomkamp et al., 2019). Similarly, *Scn1b* is a
473 gene encoding a voltage-gated sodium channel whose expression correlates with the width of the spike and
474 many other electrophysiological properties (e.g., the first principal component) for both excitatory and inhibitory
475 neuron classes. Thus, based on this data, it is perhaps expected that waveform shape would be one of the
476 most reliable electrophysiological modalities for delineating cell types *in vivo*. In addition, this close covariation
477 between transcriptomics and electrophysiological properties is perhaps the reason why our PhysMAP and previous
478 WaveMAP approaches are successful in delineating transcriptomic cell types from electrophysiological recordings.

479 When conducting this study, we also realized that a key requirement for PhysMAP is high-quality ground truth
480 datasets that cover the full range cell types. However, high-quality single unit recordings with opto-tagging are
481 difficult to obtain for at least two main reasons: opto-tagging yield/specificity and single unit isolation.

482 First, opto-tagging experiments are technically challenging due to light attenuation through brain tissue or lack
483 of transgene expression/specificity (Li et al., 2019), yield per recording is especially low when searching for less
484 common cell types such as VIP⁺ cells which compose only ~13% of inhibitory neurons (Prönneke et al., 2015).
485 These physically small cell types like VIP⁺ are often "missed" even on high-density probes. but are detected
486 on *ultra* high-density probes possibly due to increased signal quality and denser sampling (Ye & Shelton et al.
487 2023). Furthermore, most opto-tagging protocols isolate directly tagged cells by examining the first 10 ms after
488 the onset of a light pulse (Lima et al., 2009). However, Beau et al. (2024) find that some cells that fire within
489 this 10 ms window after stimulation, have their spiking abolished by the introduction of blockers of synaptic
490 transmission. Thus, more experiments are needed to refine opto-tagging experimental protocols and with devices
491 that are able to catch all cell types in an unbiased manner.

492 Second, single unit isolation is a significant challenge for electrophysiology (Buccino et al., 2022). Even with
493 modern recording technologies (Jun et al., 2017) and spike sorting methods (Pachitariu et al., 2016), single
494 neuron isolation is difficult (Vincent and Economo, 2023). Poor isolation will lead to a multi-unit being classified
495 as a single neuron and impacts PhysMAP because it blurs the physiological differences between cell types. For
496 example, averaging together waveforms from a fast-firing narrow-spiking PV⁺ cell and a slow-firing broad-spiking
497 excitatory pyramidal cell could lead to an artifactual intermediate width waveform neuron that is neither PV⁺
498 nor an excitatory pyramidal neuron. This mixing of cell types would complicate inferences from PhysMAP.

499 We anticipate that these technical challenges will be overcome. However, one key need for PhysMAP and other
500 cell type classifiers is a more exhaustive characterization of the many cell types in the brain. Inhibitory cell types
501 receive the most attention in opto-tagging experiments whereas excitatory populations have typically been under
502 explored. In the datasets we analyzed here, only one (CellExplorer data from Buzsaki lab) contains positive
503 identification of excitatory cells and this from an older study that precedes optogenetics (Henze et al., 2000).
504 This is possibly the result of reliance on classical heuristics in the field that broad-spiking cells are excitatory
505 neurons (McCormick et al., 1985). However, not all broad-spiking cells are excitatory (Wu et al., 2023) and
506 not all excitatory cells are broad-spiking (Vigneswaran et al., 2011). This complex relationship between spike
507 width and cell types is especially true in primates (Lemon et al., 2021). Another possible reason for this lack of
508 attention to direct tagging of excitatory cells is that it can be hard to distinguish which cells are being directly
509 stimulated due to the density of labeling and degree of post-synaptic secondary stimulation (Li et al., 2019). One
510 way to label excitatory neurons more sparsely and specifically, is to use projection targeting approaches identifying
511 neurons with long-range compared as opposed to local connections (Economo et al., 2018). Building full cell
512 type classifiers (for cortex) will necessitate the opto-tagging of major projection classes most notably the three
513 major excitatory types (intratelencephalic-, extratelencephalic-, and corticothalamic-projecting) and also two of
514 the major remaining major inhibitory cell types (*Lamp5* and *Sncg*). Furthermore, it should be explored the
515 extent to which electrophysiological signatures of these cell types are shared between brain areas. The conserved
516 transcriptomic identity of many cell types across areas seems to suggest that this is the case at least for cortex
517 (Yao et al., 2023).

518 Thus, PhysMAP and other cell type classifiers (Beau et al., 2024) would benefit from dedicated opto-tagging
519 experiments that assess a wide of cell types, prioritizing single unit isolation, and with robust opto-tagging
520 yields. Fortunately, two technological advances will help with these experiments. First, Neuropixels probes have
521 dramatically increased the number of recording sites and capture the same neuron across multiple channels to
522 provide a “multi-channel waveform”. This has been shown to better capture morphological details which might
523 be important for differentiating cell types that only differ in their structure (Jia et al., 2019; Ye et al., 2023;
524 Sibille et al., 2022; Carr et al., 2024). Second, new enhancer-based viruses demonstrate excellent cell subtype-
525 specificity (Green et al., 2023) along with efficacy across diverse organisms including humans (Graybuck et al.,
526 2021). The combination of these technical advances with our PhysMAP approach will likely provide robust cell
527 type classifiers that will allow the simultaneous assessment of the dynamics of neuronal cell types *in vivo* during
528 sensation, perception, and action in multiple species.

529 (Discussion: 1852 words)

530 Methods

531 The methods are organized as follows. We first describe the datasets that we used for the paper and then
532 describe the weighted nearest neighbor method from PhysMAP. Finally, we describe the various analyses used
533 for quantifying the performance of PhysMAP relative to individual modalities and also other controls.

534 *Open Datasets*

535 Here, we only briefly detail the most relevant aspects of each dataset for our purpose of validating our PhysMAP
536 approach (Table 1). We refer the reader to each dataset’s respective publication for additional methodological
537 information.

538 *Juxtacellular Mouse S1 Dataset*

539 The juxtacellular dataset used in Fig. 2 and Fig. 3 and analyzed in section were collected by Yu et al. (2019)
540 and downloaded from the associated file sharing portal (see Table 1). Recordings were performed in the primary
541 somatosensory (barrel) cortex of Sst-IRES-Cre \times Ai32, Pvalb-IRES-Cre/Pvalb-Cre \times Ai32, or Vip-IRES-Cre
542 \times Ai32 mice. These experiments focused on *in vivo* juxtacellular electrophysiological recordings using glass
543 micropipettes. After the end of recording, the cells were filled with biocytin/neurobiotin. Filled neurons both
544 acted as a verification of the recorded cell type via their morphology and also served as landmarks to align
545 recording depths of each cell with cortical layers as determined by histology.

546 ChR2-expressing neurons were located during recording by observing laser-evoked spikes that occurred 1-2 ms
547 after the onset of brief light pulses (5 ms) delivered at low frequency (5 Hz) for 3 out of every 10 s. Spikes
548 were examined by eye and non-spike events were removed after projection in PCA-space. To control for drift,
549 only waveforms between the 25th and 75th percentile of all spikes from a cell were used in further waveform
550 shape analyses (average waveform and waveform shape metrics). Averaged waveform shapes and ISI distribution
551 were extracted for each neuron from the *spkwaveform_all.mat* file in the corresponding cell type structs. In
552 addition, the waveform shape metrics of spike width and peak-to-trough ratio were also collected from this
553 file. Cells labeled “fake SOM⁺” were reassigned to the PV⁺ population and cells labeled “putative SOM⁺” were
554 reassigned to SOM⁺. This reassignment was based on a careful investigation in Yu et al. (2019) confirming earlier
555 findings that off-target recombination occurs in fast-spiking neurons of somatostatin-IRES-Cre mice possibly due
556 to transient Cre expression during development (Ma et al., 2006; Hu et al., 2013). PSTHs were calculated only
557 from whisker touches during active whisking and activity resulting from both whisker protraction and retraction
558 were used. Spiking was aligned to touch-onset and a PSTH calculated using 1 ms time bins. As not all cells
559 contained PSTHs, only the subset of units containing all three modalities (waveform shape, ISI distribution, and
560 PSTH) were used. To create the “concatenated” modality, we simply concatenated the waveform shape, ISI
561 distribution, and PSTH for a given unit into a single long feature vector. Waveform metrics—peak-to-trough
562 duration and peak-to-trough ratio—were calculated in the originating publication as the time from the first peak
563 to the first trough and the ratio of absolute values of the peak to the trough respectively.

564 *Extracellular Mouse A1 Dataset*

565 The extracellular data we analyzed from Lakunina et al. (2020) was downloaded from GitHub (see Table 1).
566 PV⁺ and SOM⁺ cells were identified by opto-tagging in Pvalb-Cre or Sst-Cre mice respectively that were the

Study	N	Available modalities	Classifiable cell types	Data URL
Yu et al. (2019)	246	Waveform shape, ISI dist., PSTH	PV ⁺ -4, PV ⁺ -5, SOM ⁺ , E-4, E-5	Link
Lakunina et al. (2020)	373	Waveform shape, ISI dist.	PV ⁺ , SOM ⁺	Link
Petersen et al. (2021)	430	Waveform shape, ISI dist., Auto-correlation, Derived ephys. metrics	PV ⁺ , SOM ⁺ , E	Link

Table 1: The three datasets analyzed with the modalities used and the cell types identified in each.

567 progeny of crossing LSL-ChR2 mice. Photoidentification of ChR2-expressing cells was relatively conservative
568 with positively identified cells being those that exhibited significant firing rate changes ($p < 0.001$) in firing rate
569 during the first 10 ms of stimulation-onset. Recordings were collected from primary auditory cortex via silicon
570 probe and stimulation was via an optical fiber located at the top of the probe recording sites. Spikes were sorted
571 offline in Klustakwik and the units used in further analyses only if ISI violation rate (number of spike intervals
572 < 2 ms) was less than 2% per cluster. In the original publication, only cells with spike quality index (SQI; ratio
573 between the peak amplitude of the waveform and the average variance, calculated using the channel with the
574 largest amplitude) above 2.5 were used; in our analysis, we only used cells with SQI above 4.0. Mean waveforms
575 were calculated for a unit using the channel with greatest amplitude of spikes.

576 *Extracellular Mouse Visual Cortex and Hippocampus Dataset (CellExplorer)*

577 The third dataset was obtained from the CellExplorer package (Petersen et al., 2021) and is heterogeneous being
578 composed of two separate extracellular electrophysiological data from the Allen Institute's Brain Observatory
579 and the Buzsaki lab and included waveform shape, ISI distribution, autocorrelogram (ACG), and various derived
580 electrophysiological metrics. We used the derived electrophysiological metrics that all units shared and this included
581 the following eleven:

- 582 • Spike width (*troughToPeak*)
- 583 • Peak-to-trough derivative (*troughtoPeakDerivative*)
- 584 • Pre-hyperpolarization peak-to-post depolarization peak ratio (*ab_ratio*)
- 585 • Coefficient of variation (*cv2*) measuring ISI distribution
- 586 • ACG tau rise (*acg_tau_rise*)
- 587 • ACG tau decay (*acg_tau_decay*)
- 588 • ACG tau bursts (*acg_tau_burst*)
- 589 • ACG refractory period (*acg_refrac*)
- 590 • ACG decay amplitude (*acg_c*)
- 591 • ACG rise amplitude (*acg_d*)
- 592 • ACG burst amplitude (*acg_h*).

593 More information on these metrics are available on the CellExplorer website docs ([link](#)). For the analyses, these
594 metrics were concatenated for each neuron into a single vector per-unit to form the electrophysiological metrics
595 modality. Given the large number of publications and protocols that have contributed to the CellExplorer dataset,
596 we only provide a brief overview here.

597 *CellExplorer: Allen Institute Data*

598 Data from the Allen Institute was collected from mouse primary visual cortex (V1) and higher visual areas (HVAs)
599 using simultaneous Neuropixels probe recordings (Siegle et al., 2021). Opto-tagging was conducted for PV⁺,
600 SOM⁺, and VIP⁺ cell types using Pvalb-IRES-Cre × Ai32, Sst-IRES-Cre × Ai32, and Vip-IRES-Cre × Ai32
601 mice respectively. Stimulation was conducted at the end of each behavior session via blue LED/laser through an
602 optical fiber with surface illumination of the cranial window using a rounded 10 ms square pulse. Tagged cells
603 were identified by an average two times increase in base firing rate during stimulation relative to baseline.

604 *CellExplorer: Buzsaki Lab Data*

605 Data from the Buzsaki lab were collected under several different protocols but CA1 excitatory neurons were
606 identified by paired intra- and extra-cellular recordings (Henze et al., 2000). V1 PV⁺ cells from PV-Cre × Ai32
607 mice were identified as those that were stimulated to over 8 standard deviations of the mean of baseline in the
608 first 6 ms of stimulation with a 10 ms square pulse (Senzai et al., 2019). Hippocampal PV⁺ cells were obtained
609 from PV-Cre × Ai32 using a 50 - 100 ms light pulse and identified as tagged if they 1) had a statistically
610 significant increase in firing rate and 2) if the stimulated firing rate was over 50% of the baseline (English &
611 McKenzie et al. 2017).

612 *PhysMAP Preprocessing*

613 Modalities from all neurons were preprocessed via the same steps required by and performed using Seurat v4:

- 614 1. **PCA reduction:** All modalities were initially dimensionality reduced to the smallest input modality di-
615 mension. For the S1 dataset, this is 30-dimensions; for the A1 dataset, this is 20-dimensions; and for
616 CellExplorer, this is 40-dimensions.
- 617 2. **Normalization:** Each modality is separately normalized using a centered log ratio transform.
- 618 3. **Rescaling and centering:** Each feature of each modality is linearly rescaled to occupy the same unit
619 variance. Each feature is then centered by mean subtraction.

620 *Weighted Nearest Neighbor (WNN) Algorithm*

621 The PhysMAP approach uses the WNN algorithm available from Seurat v4 (Hao & Hao et al. 2021) which is
622 summarized as the following steps.

- 623 1. **Nearest neighbor identification:** Within the preprocessed, high-dimensional space of a given modality
624 (here, waveform shape Fig. 1A), a single unit is selected (Fig. 1B, blue sphere) and its nearest neighbors
625 identified via Euclidean distance in the ambient (un-dimensionality reduced) space (Fig. 1B, red spheres).
- 626 2. **Nearest neighbor prediction (within modality):** A prediction of the selected single unit's waveform is
627 made via unweighted averaging the k -nearest neighbors identified in Fig. 1B. By default, we use $k = 20$
628 and Euclidean distances in the ambient space to locate the nearest neighbors.
- 629 3. **Within-modal affinity calculation:** The unit's nearest neighbor predicted waveform shape generated by
630 k -nearest neighbors ($\hat{x}_{i,knn=20}$) and its actual waveform shape (x_i) are then passed into a modified UMAP
631 distance kernel,

$$d_{\text{UMAP}}(x_i, \hat{x}_i) \equiv \exp \frac{-\max\left(0, d_{L^2}(x_i, \hat{x}_{i,knn=20}) - d_{L^2}(x_i, x_{i,knn=1})\right)}{\sigma_i - d_{L^2}(x_i, x_{i,knn=1})}$$

629 where $d_{L^2}(\cdot)$ is the Euclidean (L^2 -norm) distance metric; $d_{L^2}(x_i, x_{i,knn=1})$ is the Euclidean distance from
630 x_i to its nearest neighbor; and σ_i is a modified UMAP bandwidth equal to the average of the Euclidean
631 distances from the i^{th} unit to the 20 nearest units with lowest non-zero Jaccard index. This equation is
632 used to calculate the “within-modal affinity” which gives a measure of how predictive a certain modality is
633 of this unit. This forms the numerator of the “affinity ratio” shown in Fig. 1D.

- 634 4. **Nearest neighbor identification (cross modality):** Examining normalized single unit ISI distributions
635 (the “cross-modal” space) generated from the same neurons in Fig. 1E, we locate the same unit and
636 nearest neighbors previously identified but in ISI distribution-space Fig. 1F.
- 637 5. **Nearest neighbor prediction (cross modality):** As in Fig. 1C, a nearest neighbor prediction is made
638 but this time, in a different modality albeit using the same neurons identified in the original modality Fig.
639 1G, red dashed circle.
- 640 6. **Affinity ratio calculation:** This unit and its prediction in the cross modality (ISI distribution-space)
641 are used to compute the “cross-modal affinity” by passing it into a modified UMAP distance kernel; this
642 forms the denominator of the waveform affinity ratio for a given unit ($S_{\text{Wave}}(i)$) shown in Fig. 1D. This
643 process is repeated for every unit in the dataset and also in the reverse manner: beginning with the ISI
644 distribution-space as the within modality and with the waveform shape-space as the cross modality. For
645 datasets with more than two modalities, the affinity ratio is calculated for each modality versus every
646 other and for every unit. Thus for a dataset with d data points and n feature modalities, there will be
647 $d \cdot 2 \frac{n!}{2!(n-2)!} = d \cdot n(n-1)$ affinity ratios to calculate.
- 648 7. **Converting affinity ratios to modality weights:** With affinity ratios calculated for both waveform and
649 ISI distribution across all units, modality weights ($\beta_{\text{Wave}}(i)$ and $\beta_{\text{ISI}}(i)$) for each unit can be calculated.
650 For a given unit, this weight is the ratio of a given modality's affinity ratio for that unit divided by the
651 sum of all other modality's affinity ratios for that unit. For the two-modality case, this is summarized by
652 the equation in Fig. 1H.

- 653 8. **Calculating unit pair-wise affinities:** For every pair of units, an affinity is also calculated by passing both
654 of them and their coordinates into UMAP's distance kernel. This occurs for the same pair of units in both
655 the waveform shape-space Fig. 1I, top and the ISI distribution-space Fig. 1I, bottom to determine pair-wise
656 affinities $\theta_{\text{Wave}}(i, j)$ and $\theta_{\text{ISI}}(i, j)$ respectively. This is done for all pairs of points and all modalities.
- 657 9. **Creating the WNN:** To create the weighted-nearest neighbors representation, a modality-weighted sum
658 of the pair-wise affinities Fig. 1J, top is taken to produce a "connectivity matrix" Fig. 1J, bottom. The k -
659 nearest neighbor algorithm is applied to this matrix to form the final WNN (default number of neighbors is
660 $k = 200$) which is then visualized into two dimensions with UMAP's force-directed graph layout projection
661 for visualization.

662 *Classification Analysis*

663 After a WNN is constructed for a dataset, we used UMAP's force-directed graph layout procedure to project
664 the graph into a 10-dimensional embedded space. We then used a random 15-85% test-train split and trained
665 a stochastic gradient boosted tree model (GBM) classifier on the training set with five-fold cross-validation
666 using the *caret* package in R (Kuhn, 2008) to identify each underlying cell type with multiclass (one-vs-rest)
667 objective function. We then repeated this train-test split 20 times and averaged the balanced accuracy over
668 these independent runs. Each classifier also underwent hyperparameter tuning via a grid search. This procedure
669 was run identically for the controls in Fig. S2A, B except with varying graph embedding dimension and using
670 different classifiers available in the *caret* package.

671 *Leiden Community Detection*

To calculate the MARI score for each dataset across different numbers of clusters we used Leiden clustering (Traag et al., 2019) on each dataset's WNN graphs with resolution between 0.1 and 3.0 in 0.1 resolution intervals. The Leiden algorithm is a method for "community detection" which finds highly inter-connected nodes on a network graph akin to clustering in a metric space. Algorithms like Leiden (and the simpler Louvain algorithm) attempt to find a partitioning of the network into a set of communities that maximizes the modularity of each community. This modularity, \mathcal{H} , is a measure of how inter-connected the nodes are within a community versus outside of it. In the below definition, m is the total number of edges in the network, e_c is the number of edges in a community c , γ is the resolution parameter, and K_c is the sum of the degrees of the nodes in community c .

$$\mathcal{H} \equiv \frac{1}{2m} \sum_c \left(e_c - \gamma \frac{K_c^2}{2m} \right)$$

672 Thus, the resolution parameter is effectively a prior on the number of expected communities that should be
673 found during the optimization; with lower resolution, less communities are found and with higher resolution,
674 more communities are found.

675 *Modified Adjusted Rand Index (MARI) Calculation*

676 Once this partitioning of a graph into communities is computed, a MARI score (Sundqvist et al., 2020) is
677 calculated between each of these each neuron's community membership and the ground truth cluster identity to
678 determine how closely they corresponded. To understand MARI, which is a modification of the adjusted Rand
679 index (ARI), we begin with an explanation of the Rand index. Given two clusterings, X and Y of the same
680 dataset, the number of pair-wise elements that share cluster membership in both X and Y is a and the number
681 of pair-wise elements that share different cluster memberships is b . The Rand index (RI) is $RI = \frac{a+b}{\binom{n}{2}}$ where n is
682 the total number of pairs. The Rand index can also be interpreted as the sum of the number of true positive and
683 true negative correspondences divided by the total number of guesses if one clustering is regarded as a classifier
684 prediction of the other. However, in the classifier interpretation, the Rand index is not corrected for chance
685 predictions. The ARI corrects for these chance predictions by subtracting the expected Rand index under the
686 hypergeometric distribution given two independent clusterings and fixed number of clusters in each. MARI further

687 refines this metric by instead incorporating a multinomial distribution which does not enforce cluster size and is
688 a better assumption given that most clustering algorithms (including Leiden) do not fix cluster sizes. This was
689 used to compare each of the Leiden clusterings at different resolutions to underlying ground truth clusters.

690 *Reference Mapping with Anchors and Cell Type Classification*

Unlike the previous classification analyses, an 80-20 train-test split was conducted before applying PhysMAP. This was done to combat bias in cross-validation because of data leakage due to pre-processing transformations (Moscovich and Rosset, 2022). The training set was used to create the reference mapping upon and the test set was used to create the query dataset. To construct the reference mapping, an WNN was constructed from each neuron's waveform shape, ISI distribution, autocorrelogram, and derived electrophysiological metrics as before. Following the procedure for multi-modal reference mapping recommended in Hao & Hao et al. 2021, we first ran supervised principal component analysis (SPCA) on the reference data. Next, the query dataset is projected onto the reference using the previously computed SPCA transform. Now that reference and query datasets occupy the same space, "anchors" can be calculated between the reference and query. These anchors are pairs of cells between the reference and query that are located within each other's neighborhoods; this concept is also referred to as mutual nearest neighbors (MNN). These neighborhoods are defined by computing k -nearest neighbors with $k = 5$. This nearest neighbor search is conducted in only the top-30 SPC's. Once all anchors are found, they are scored which is an assessment of how confident we are in their correspondences. This score is calculated for each anchor given the 30 nearest neighbors in each reference and query dataset for both the reference and query data points of the anchor respectively. The overlap of these nearest neighbor matrices between the 0.01 and 0.90 quantiles is then linearly rescaled to be between 0 and 1 for all reference-query anchor pairs; this provides a score for each anchor. A weight matrix is then constructed between each query cell c and each i th nearest anchor a_i (where $i \in [1, 50]$) based upon the distance to each query cell and the corresponding anchor score. These weighted distances $D_{c,i}$ are calculated as

$$D_{c,i} = S_{a_i} \left(1 - \frac{\|c, a_i\|}{\|c, a_{50}\|} \right)$$

where S_{a_i} is the score of the i th anchor where $\|\cdot\|$ is the Euclidean distance. These distances are passed through a Gaussian distance kernel $\tilde{D}_{c,i}$ as,

$$\tilde{D}_{c,i} = 1 - \frac{-D_{c,i}}{e^4}$$

to form the entries of a weight matrix

$$W_{c,i} = \frac{\tilde{D}_{c,i}}{\sum_{j=50} \tilde{D}_{c,i}}$$

For cell type classification, a classification matrix L is created where each row corresponds to a ground truth cell type class and each column corresponds to each reference anchor. If a certain reference anchor pertains to a certain cell type, it is given a 1 and a 0 if otherwise. Label predictions P are then computed simply as,

$$P = LW^T.$$

691 This returns a prediction score for each cell in the query dataset. A query cell is thus given a predicted cell type
692 of whichever class has the highest prediction score.

References

- Adesnik H, Naka A (2018). *Cracking the function of layers in the sensory cortex*. *Neuron*, 100(5):1028–1043
- Aharoni D, Hoogland TM (2019). *Circuit Investigations With Open-Source Miniaturized Microscopes: Past, Present and Future*. *Frontiers in Cellular Neuroscience*, 13:141. ISSN 1662-5102
- Amatrudo JM, Weaver CM, Crimins JL, Hof PR, Rosene DL, Luebke JI (2012). *Influence of highly distinctive structural properties on the excitability of pyramidal neurons in monkey visual and prefrontal cortices*. *Journal of Neuroscience*, 32(40):13644–13660
- Argelaguet R, Arnol D, Bredikhin D, Deloro Y, Velten B, Marioni JC, Stegle O (2020). *MOFA+: a statistical framework for comprehensive integration of multi-modal single-cell data*. *Genome biology*, 21(1):1–17
- Bakken TE, Jorstad NL, Hu Q, Lake BB, Tian W, Kalmbach BE, Crow M, et al. (2021). *Comparative cellular analysis of motor cortex in human, marmoset and mouse*. *Nature*, 598(7879):111–119
- Beau M, Herzfeld DJ, Naveros F, Hemelt ME, D'Agostino F, Oostland M, Sánchez-López A, et al. (2024). *A deep-learning strategy to identify cell types across species from high-density extracellular recordings*. bioRxiv (BICCN) BICCN (2021). *A multimodal cell census and atlas of the mammalian primary motor cortex*. *Nature*, 598(7879):86–102
- Bliss-Moreau E, Costa VD, Baxter MG (2022). *A pragmatic reevaluation of the efficacy of nonhuman primate optogenetics for psychiatry*. *Oxford Open Neuroscience*, 1:kvac006
- Bloch J, Greaves-Tunnell A, Shea-Brown E, Harchaoui Z, Shojaie A, Yazdan-Shahmorad A (2022). *Network structure mediates functional reorganization induced by optogenetic stimulation of non-human primate sensorimotor cortex*. *iScience*, 25(5):104285. ISSN 2589-0042
- Bohlen M, Tremblay S (2023). *Vectorology for Optogenetics and Chemogenetics*. *NeuroMethods*, pages 293–329. ISSN 0893-2336
- Bomkamp C, Tripathy SJ, Gonzales CB, Hjerling-Leffler J, Craig AM, Pavlidis P (2019). *Transcriptomic correlates of electrophysiological and morphological diversity within and across excitatory and inhibitory neuron classes*. *PLoS Computational Biology*, 15(6):e1007113. ISSN 1553-734X
- Boucher PO, Wang T, Carceroni L, Kane G, Shenoy KV, Chandrasekaran C (2023). *Initial conditions combine with sensory evidence to induce decision-related dynamics in premotor cortex*. *Nature Communications*, 14(1):6510
- Buccino AP, Garcia S, Yger P (2022). *Spike sorting: new trends and challenges of the era of high-density probes*. *Progress in Biomedical Engineering*, 4(2):022005
- Carr N, Zhu S, Lee K, Xia R, Chen X, Perliss A, Moore T, Chandrasekaran C (2024). *Neuropixels reveal microcircuits in monkey primary visual cortex in-vivo*. bioRxiv
- Cembrowski MS, Menon V (2018). *Continuous variation within cell types of the nervous system*. *Trends in Neurosciences*, 41(6):337–348
- Chen S, Liu Y, Wang ZA, Colonell J, Liu LD, Hou H, Tien NW, et al. (2024). *Brain-wide neural activity underlying memory-guided movement*. *Cell*, 187(3):676–691.e16. ISSN 0092-8674
- Cheng M, Li Z, Costa IG (2022). *MOJITOO: a fast and universal method for integration of multimodal single-cell data*. *Bioinformatics*, 38(Supplement_1):i282–i289
- Christensen AJ, Ott T, Kepecs A (2022). *Cognition and the single neuron: How cell types construct the dynamic computations of frontal cortex*. *Current Opinion in Neurobiology*, 77:102630
- Dudok B, Szoboszlay M, Paul A, Klein PM, Liao Z, Hwaun E, Szabo GG, et al. (2021). *Recruitment and inhibitory action of hippocampal axo-axonic cells during behavior*. *Neuron*, 109(23):3838–3850.e8. ISSN 0896-6273
- Economo MN, Viswanathan S, Tasic B, Bas E, Winnubst J, Menon V, Graybiel LT, et al. (2018). *Distinct descending motor cortex pathways and their roles in movement*. *Nature*, 563(7729):79–84. ISSN 0028-0836
- English DF, McKenzie S, Evans T, Kim K, Yoon E, Buzsáki G (2017). *Pyramidal Cell-Interneuron Circuit Architecture and Dynamics in Hippocampal Networks*. *Neuron*, 96(2):505–520.e7. ISSN 0896-6273
- Esparza J, Sebastián ER, Prida LMdl (2023). *From cell types to population dynamics: Making hippocampal manifolds physiologically interpretable*. *Current Opinion in Neurobiology*, 83:102800. ISSN 0959-4388
- Ferguson KA, Cardin JA (2020). *Mechanisms underlying gain modulation in the cortex*. *Nature Reviews Neuroscience*, 21(2):80–92. ISSN 1471-003X

- Gouwens NW, Sorensen SA, Baftizadeh F, Budzillo A, Lee BR, Jarsky T, Alfiler L, et al. (2020). *Integrated morphoelectric and transcriptomic classification of cortical GABAergic cells*. *Cell*, 183(4):935–953
- Graybuck LT, Daigle TL, Sedeño-Cortés AE, Walker M, Kalmbach B, Lenz GH, Morin E, et al. (2021). *Enhancer viruses for combinatorial cell-subclass-specific labeling*. *Neuron*, 109(9):1449–1464.e13. ISSN 0896-6273
- Green J, Bruno CA, Traunmüller L, Ding J, Hrvatin S, Wilson DE, Khodadad T, Samuels J, Greenberg ME, Harvey CD (2023). *A cell-type-specific error-correction signal in the posterior parietal cortex*. *Nature*, 620(7973):366–373. ISSN 0028-0836
- Haghverdi L, Lun ATL, Morgan MD, Marioni JC (2018). *Batch effects in single-cell RNA-sequencing data are corrected by matching mutual nearest neighbors*. *Nature Biotechnology*, 36(5):421–427. ISSN 1087-0156
- Hao Y, Hao S, Andersen-Nissen E, Mauck WM, Zheng S, Butler A, Lee MJ, et al. (2021). *Integrated analysis of multimodal single-cell data*. *Cell*, 184(13):3573–3587.e29. ISSN 0092-8674
- Hao Y, Stuart T, Kowalski MH, Choudhary S, Hoffman P, Hartman A, Srivastava A, et al. (2023). *Dictionary learning for integrative, multimodal and scalable single-cell analysis*. *Nature Biotechnology*, pages 1–12. ISSN 1087-0156
- Henze DA, Borhegyi Z, Csicsvari J, Mamiya A, Harris KD, Buzsáki G (2000). *Intracellular Features Predicted by Extracellular Recordings in the Hippocampus In Vivo*. *Journal of Neurophysiology*, 84(1):390–400. ISSN 0022-3077
- Heynen AJ, Bear MF (2001). *Long-Term Potentiation of Thalamocortical Transmission in the Adult Visual Cortex In Vivo*. *The Journal of Neuroscience*, 21(24):9801–9813. ISSN 0270-6474
- Hostetler RE, Hu H, Agmon A (2023). *Genetically Defined Subtypes of Somatostatin-Containing Cortical Interneurons*. *eNeuro*, 10(8)
- Hu H, Cavendish JZ, Agmon A (2013). *Not all that glitters is gold: off-target recombination in the somatostatin-IRES-Cre mouse line labels a subset of fast-spiking interneurons*. *Frontiers in Neural Circuits*, 7:195. ISSN 1662-5110
- Hussar CR, Pasternak T (2009). *Flexibility of sensory representations in prefrontal cortex depends on cell type*. *Neuron*, 64(5):730–743
- Jia X, Siegle JH, Bennett C, Gale SD, Denman DJ, Koch C, Olsen SR (2019). *High-density extracellular probes reveal dendritic backpropagation and facilitate neuron classification*. *Journal of Neurophysiology*, 121(5):1831–1847. ISSN 0022-3077
- Jun JJ, Steinmetz NA, Siegle JH, Denman DJ, Bauza M, Barbarits B, Lee AK, et al. (2017). *Fully integrated silicon probes for high-density recording of neural activity*. *Nature*, 551(7679):nature24636. ISSN 1476-4687
- Keller AJ, Dipoppa M, Roth MM, Caudill MS, Ingrosso A, Miller KD, Scanziani M (2020). *A Disinhibitory Circuit for Contextual Modulation in Primary Visual Cortex*. *Neuron*. ISSN 0896-6273
- Kim D, Jeong H, Lee J, Ghim JW, Her ES, Lee SH, Jung MW (2016). *Distinct Roles of Parvalbumin- and Somatostatin-Expressing Interneurons in Working Memory*. *Neuron*, 92(4):902–915. ISSN 0896-6273
- Kriegeskorte N, Simmons WK, Bellgowan PSF, Baker CI (2009). *Circular analysis in systems neuroscience: the dangers of double dipping*. *Nature Neuroscience*, 12(5):nn.2303. ISSN 1097-6256
- Krimer LS, Zaitsev AV, Czanner G, Kroner S, González-Burgos G, Povysheva NV, Iyengar S, Barrionuevo G, Lewis DA (2005). *Cluster analysis-based physiological classification and morphological properties of inhibitory neurons in layers 2–3 of monkey dorsolateral prefrontal cortex*. *Journal of neurophysiology*, 94(5):3009–3022
- Kuhn M (2008). *Building Predictive Models in R Using the caret Package*. *Journal of Statistical Software*, 28(5):1–26
- Kullander K, Topolnik L (2021). *Cortical disinhibitory circuits: cell types, connectivity and function*. *Trends in Neurosciences*, 44(8):643–657
- Kvitsiani D, Ranade S, Hangya B, Taniguchi H, Huang JZ, Kepecs A (2013). *Distinct behavioural and network correlates of two interneuron types in prefrontal cortex*. *Nature*, 498(7454):363–366. ISSN 0028-0836
- Lakunina AA, Nardoci MB, Ahmadian Y, Jaramillo S (2020). *Somatostatin-Expressing Interneurons in the Auditory Cortex Mediate Sustained Suppression by Spectral Surround*. *The Journal of Neuroscience*, 40(18):3564–3575. ISSN 0270-6474

- Latuske P, Toader O, Allen K (2015). *Interspike intervals reveal functionally distinct cell populations in the medial entorhinal cortex*. *Journal of Neuroscience*, 35(31):10963–10976
- Lee EK, Balasubramanian H, Tsolias A, Anakwe SU, Medalla M, Shenoy KV, Chandrasekaran C (2021). *Non-linear dimensionality reduction on extracellular waveforms reveals cell type diversity in premotor cortex*. *eLife*, 10:e67490
- Lee K, Carr N, Perliss A, Chandrasekaran C (2023). *WaveMAP for identifying putative cell types from in vivo electrophysiology*. *STAR Protocols*, 4(2):102320. ISSN 2666-1667
- Lemon RN, Baker SN, Kraskov A (2021). *Classification of Cortical Neurons by Spike Shape and the Identification of Pyramidal Neurons*. *Cerebral Cortex*, pages bhab147–. ISSN 1047-3211
- Li N, Chen S, Guo ZV, Chen H, Huo Y, Inagaki HK, Chen G, et al. (2019). *Spatiotemporal constraints on optogenetic inactivation in cortical circuits*. *eLife*, 8:e48622
- Lima SQ, Hromádka T, Znamenskiy P, Zador AM (2009). *PINP: A New Method of Tagging Neuronal Populations for Identification during In Vivo Electrophysiological Recording*. *PLoS ONE*, 4(7):e6099
- Ma Y, Hu H, Berrebi AS, Mathers PH, Agmon A (2006). *Distinct Subtypes of Somatostatin-Containing Neocortical Interneurons Revealed in Transgenic Mice*. *The Journal of Neuroscience*, 26(19):5069–5082. ISSN 0270-6474
- McCormick DA, Connors BW, Lighthall JW, Prince DA (1985). *Comparative electrophysiology of pyramidal and sparsely spiny stellate neurons of the neocortex*. *Journal of Neurophysiology*, 54(4):782–806. ISSN 0022-3077
- McInnes L, Healy J, Melville J (2018). *UMAP: Uniform Manifold Approximation and Projection for Dimension Reduction*. *arXiv*
- Millman DJ, Ocker GK, Caldejon S, Kato I, Larkin JD, Lee EK, Luviano J, et al. (2020). *VIP interneurons in mouse primary visual cortex selectively enhance responses to weak but specific stimuli*. *eLife*, 9:e55130
- Mitchell JF, Sundberg KA, Reynolds JH (2007). *Differential attention-dependent response modulation across cell classes in macaque visual area V4*. *Neuron*, 55(1):131–141
- Moscovich A, Rosset S (2022). *On the Cross-Validation Bias due to Unsupervised Preprocessing*. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 84(4):1474–1502. ISSN 1369-7412
- Onorato I, Tzanou A, Schneider M, Uran C, Broggin A, Vinck M (2023). *Distinct roles of PV and Sst interneurons in visually-induced gamma oscillations*. *bioRxiv*
- O'Shea DJ, Trautmann E, Chandrasekaran C, Stavisky S, Kao JC, Sahani M, Ryu S, Deisseroth K, Shenoy KV (2017). *The need for calcium imaging in nonhuman primates: New motor neuroscience and brain-machine interfaces*. *Experimental Neurology*, 287:437–451. ISSN 0014-4886
- Pachitariu M, Steinmetz N, Kadir S, Carandini M, D HK (2016). *Kilosort: realtime spike-sorting for extracellular electrophysiology with hundreds of channels*. *bioRxiv*, page 061481
- Peixoto D, Verhein JR, Kiani R, Kao JC, Nuyujukian P, Chandrasekaran C, Brown J, et al. (2021). *Decoding and perturbing decision states in real time*. *Nature*, pages 1–6. ISSN 0028-0836
- Petersen PC, Hernandez M, Buzsáki G (2020). *The Buzsáki Lab Databank - Public electrophysiological datasets from awake animals*
- Petersen PC, Siegle JH, Steinmetz NA, Mahallati S, Buzsáki G (2021). *CellExplorer: A framework for visualizing and characterizing single neurons*. *Neuron*, 109(22):3594–3608.e2. ISSN 0896-6273
- Pinto L, Dan Y (2015). *Cell-Type-Specific Activity in Prefrontal Cortex during Goal-Directed Behavior*. *Neuron*, 87(2):437–450. ISSN 0896-6273
- Povyshva NV, Zaitsev AV, Gonzalez-Burgos G, Lewis DA (2013). *Electrophysiological Heterogeneity of Fast-Spiking Interneurons: Chandelier versus Basket Cells*. *PLOS One*, 8(8):1–13
- Prönneke A, Scheuer B, Wagener RJ, Möck M, Witte M, Staiger JF (2015). *Characterizing VIP Neurons in the Barrel Cortex of VIPcre/tdTomato Mice Reveals Layer-Specific Differences*. *Cerebral Cortex*, 25(12):4854–4868. ISSN 1047-3211
- Ramadan M, Lee EK, Vries Sd, Caldejon S, Kato I, Roll K, Griffin F, et al. (2022). *A Standardized Nonvisual Behavioral Event Is Broadcasted Homogeneously across Cortical Visual Areas without Modulating Visual Responses*. *eNeuro*, 9(5):ENEURO.0491–21.2022

- Scala F, Kobak D, Bernabucci M, Bernaerts Y, Cadwell CR, Castro JR, Hartmanis L, et al. (2021). *Phenotypic variation of transcriptomic cell types in mouse motor cortex*. *Nature*, 598(7879):144–150
- Schneider A, Azabou M, McDougall-Vigier L, Parks DF, Ensley S, Bhaskaran-Nair K, Nowakowski T, Dyer EL, Hengen KB (2023). *Transcriptomic cell type structures in vivo neuronal activity across multiple timescales*. *Cell Reports*, 42(4):112318. ISSN 2211-1247
- Senzai Y, Fernandez-Ruiz A, Buzsáki G (2019). *Layer-Specific Physiological Features and Interlaminar Interactions in the Primary Visual Cortex of the Mouse*. *Neuron*, 101(3):500–513.e5. ISSN 0896-6273
- Sibille J, Gehr C, Benichov JI, Balasubramanian H, Teh KL, Lupashina T, Vallentin D, Kremkow J (2022). *High-density electrode recordings reveal strong and specific connections between retinal ganglion cells and midbrain neurons*. *Nature Communications*, 13(1):5218
- Siegle JH, Jia X, Durand S, Gale S, Bennett C, Graddis N, Heller G, et al. (2021). *Survey of spiking in the mouse visual system reveals functional hierarchy*. *Nature*, pages 1–7. ISSN 0028-0836
- Someck S, Levi A, Sloin HE, Spivak L, Gattegno R, Stark E (2023). *Positive and biphasic extracellular waveforms correspond to return currents and axonal spikes*. *Communications Biology*, 6(1):950
- Steinmetz NA, Zatka-Haas P, Carandini M, Harris KD (2019). *Distributed coding of choice, action and engagement across the mouse brain*. *Nature*, pages 1–8. ISSN 0028-0836
- Stuart T, Butler A, Hoffman P, Hafemeister C, Papalexi E, Mauck WM, Hao Y, Stoeckius M, Smibert P, Satija R (2019). *Comprehensive Integration of Single-Cell Data*. *Cell*, 177(7):1888–1902.e21. ISSN 0092-8674
- Sundqvist M, Chiquet J, Rigai G (2020). *Adjusting the adjusted Rand Index – A multinomial story*. *arXiv: Methodology*
- Takasaki K, Abbasi-Asl R, Waters J (2020). *Superficial bound of the depth limit of 2-photon imaging in mouse brain*. *eNeuro*, 7(1):ENEURO.0255–19.2019. ISSN 2373-2822
- Tasic B, Yao Z, Graybiuck LT, Smith KA, Nguyen TN, Bertagnolli D, Goldy J, et al. (2018). *Shared and distinct transcriptomic cell types across neocortical areas*. *Nature*, 563(7729):72–78. ISSN 0028-0836
- Traag VA, Waltman L, Eck NJ (2019). *From Louvain to Leiden: guaranteeing well-connected communities*. *Scientific Reports*, 9(1):5233
- Trautmann EM, O’Shea DJ, Sun X, Marshel JH, Crow A, Hsueh B, Vesuna S, et al. (2021). *Dendritic calcium signals in rhesus macaque motor cortex drive an optical brain-computer interface*. *Nature Communications*, 12(1):3689
- Tremblay R, Lee S, Rudy B (2016). *GABAergic Interneurons in the Neocortex: From Cellular Properties to Circuits*. *Neuron*, 91(2):260–292. ISSN 0896-6273
- Urban-Ciecko J, Barth AL (2016). *Somatostatin-expressing neurons in cortical networks*. *Nature Reviews Neuroscience*, 17(7):401–409
- Vigneswaran G, Kraskov A, Lemon RN (2011). *Large Identified Pyramidal Cells in Macaque Motor and Pre-motor Cortex Exhibit “Thin Spikes”: Implications for Cell Type Classification*. *The Journal of Neuroscience*, 31(40):14235–14242. ISSN 0270-6474
- Vincent JP, Economo MN (2023). *Assessing cross-contamination in spike-sorted electrophysiology data*. *bioRxiv*
- Wei Y, Nandi A, Jia X, Siegle JH, Denman D, Lee SY, Buchin A, et al. (2023). *Associations between in vitro, in vivo and in silico cell classes in mouse primary visual cortex*. *Nature Communications*, 14(1):2344
- Wu SJ, Sevier E, Dwivedi D, Saldi GA, Hairston A, Yu S, Abbott L, et al. (2023). *Cortical somatostatin interneuron subtypes form cell-type-specific circuits*. *Neuron*. ISSN 0896-6273
- Yao Z, Velthoven CTJv, Kunst M, Zhang M, McMillen D, Lee C, Jung W, et al. (2023). *A high-resolution transcriptomic and spatial atlas of cell types in the whole mouse brain*. *Nature*, 624(7991):317–332. ISSN 0028-0836
- Yavorska I, Wehr M (2016). *Somatostatin-Expressing Inhibitory Interneurons in Cortical Circuits*. *Frontiers in Neural Circuits*, 10:76
- Ye Z, Shelton AM, Shaker JR, Boussard JM, Colonell J, Minavi S, Chen S, et al. (2023). *Ultra-high density electrodes improve detection, yield, and cell type specificity of brain recordings*. *bioRxiv*
- Yu J, Hu H, Agmon A, Svoboda K (2019). *Recruitment of GABAergic Interneurons in the Barrel Cortex during Active Tactile Behavior*. *Neuron*, 104(2):412–427.e4. ISSN 0896-6273

Zaitsev AV, Povysheva NV, Gonzalez-Burgos G, Rotaru D, Fish KN, Krimer LS, Lewis DA (2008). *Interneuron Diversity in Layers 2–3 of Monkey Prefrontal Cortex*. *Cerebral Cortex*, 19(7):1597–1615. ISSN 1047-3211

Author Contributions

EKL and CC jointly developed the idea of a cell type classifier and ultimately a lookup table based on ideas from Dr. Karel Svoboda, and Dr. Josh Siegle. AG helped with derived waveform metric analysis and several control analyses. GH helped collect one of the open datasets and provided technical advice especially with regards to the usage of quality metric and helped edit initial manuscript drafts. AL and SJ provided technical support and also collected one of the open datasets. PFP provided guidance on various approaches for multi-modal integration and other technical advice. Both EKL and CC curated datasets. CC wrote initial R code and performed analyses. Code was further refined and augmented by EKL. EKL made figures and wrote initial drafts of the paper with CC.

Acknowledgments

We are grateful to Dr. Karel Svoboda, Dr. Anna Lakunina, and Dr. Josh Siegle for the initial ideas of cell type classifiers and lookup tables for electrophysiology. We also thank Dr. Yujin Han, Dr. Josh Siegle, Pierre Boucher, Nicole Carr, Tian Wang, Vivan Moosmann, Tushar Arora, and Munib Hasnain for their thoughtful comments on our manuscript.

Declaration of Interests

The authors declare no competing interests.